

SPACE COMPUTER CORPORATION

2800 OLYMPIC BLVD., SUITE 104 • SANTA MONICA, CALIFORNIA 90404-4119 • FAX: (213)829-1694

(213)829-7733

SCC-R-121-2

Optimum Clutter Rejection for Infrared Surveillance Sensors


A. D. Stocker

19 December 1989

Submitted to:

Naval Air Development Center
Warminster, PA 18974-5000
Attention: Mr. Stephen Campana, Code 501

Final Report under Contract N62269-89-C-0519



19950120 070

DTIC QUALITY CONTROL

SPACE COMPUTER CORPORATION
2800 Olympic Blvd., Suite 104
Santa Monica, CA 90404-4119




Table of Contents

<u>Section</u>	<u>Page</u>
1.0 Introduction and Summary	1
2.0 Results of Phase I Work	3
2.1 Statement of Problem	3
2.1.1 Existing Approaches	3
2.1.2 Fully-Adaptive Spatial Filtering	4
2.2 Phase I Technical Objectives	5
2.3 Technical Approach	5
2.4 Baseline Spatial Filter Bank	6
2.4.1 Baseline Processor Description	6
2.4.2 Application to IRST Point Target Detection	6
2.5 Adaptive Spatial Filtering	11
2.5.1 Mathematical Formulation	11
2.5.2 Properties of the Adaptive Filter	17
2.5.3 Application to IRST Point Target Detection	20
2.5.4 Theoretical Performance Prediction	25
2.6 Spatial Prefiltering	28
2.6.1 Linear High-Pass Filtering	29
2.6.2 Nonlinear High-Pass Filtering	37
2.7 Adaptive Filter Selection	42
2.8 Algorithm Performance Evaluation	43
2.8.1 TSCORE Evaluation Methodology	43
2.8.2 TSCORE Results	47
2.9 Computation Load Evaluation	57
3.0 Conclusions of Phase I Study	58
4.0 References	59

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By <i>per letter</i>	
Distribution	
Availability Codes	
Dist	Avail and/or Special
<i>A-1</i>	

1.0 Introduction and Summary

This document is the Final Report prepared under contract N62269-89-C-0519 with the Naval Air Development Center in Warminster, Pennsylvania. The contract was awarded in response to a Phase I SBIR proposal originally submitted by Space Computer Corporation in January 1988 to the Office of Naval Technology under SBIR Topic Number N88-15. The period of performance was 28 March 1989 to 28 September 1989.

In accordance with NADC's direction, the basic problem addressed in the Phase I study was the use of adaptive spatial filtering concepts to enhance the detectability of long-range targets viewed against background clutter by an airborne Infrared Search and Track (IRST) system. For the most part, these stressing threats appear to the IRST sensor as unresolved "point" targets moving against a structured background scene consisting of clouds, blue sky and possibly terrain. The IRST signal processor must be capable of detecting such targets at the lowest possible signal-to-clutter ratios, while strictly controlling the number of false alarms that must be dealt with by higher levels of data processing (e.g., scan-to-scan tracking). The detection signal processing must typically be performed on tens of thousands of pixels in two-dimensional "frames" collected at sensor scan periods on the order of 10 seconds.

Adaptive spatial filtering is a robust approach to spatial discrimination that is based on the general theory of statistical hypothesis testing. In this approach, a likelihood ratio decision statistic is used to test for the presence of a hypothesized target signal at each pixel location. Application of the test over an entire scan frame constitutes a nonlinear spatial filter whose parameters are continuously adjusted, as a function of the locally measured background clutter covariance, to optimize detection performance for a specified target response.

The potential advantages of a fully-adaptive filter with respect to existing IRST processing schemes are improved detection sensitivity in unknown, nonhomogeneous clutter backgrounds, and better false alarm regulation. The main disadvantage is the additional computational cost of implementing such a filter, which results primarily from the need to estimate and invert a clutter covariance matrix at each pixel location in the frame. This problem was clearly recognized in our Phase I proposal, where we noted that applying relatively simple prefilters to IR scenes often served to approximately "diagonalize" (or locally whiten) the background covariance and greatly simplify the adaptive filter implementation. The use of linear and nonlinear prefilters for this purpose was therefore an important topic of investigation.

Our technical approach to the Phase I study was to formulate and implement adaptive spatial filters for IRST point target detection, and to compare their performance and computational costs with those of a baseline IRST processor which employs a precomputed filter bank. Following some initial trade studies, two candidate adaptive filter configurations were selected for further evaluation. The first scheme employed linear high-pass prefiltering followed by a fully-adaptive spatial filter, with no prior assumptions about the structure of the background covariance. The second was a simplified adaptive filter designed to exploit the clutter decorrelation provided by a small-kernel median-removal filter. It was based on the assumption that the clutter covariance matrix was diagonalized by the prefiltering.

Algorithm performance evaluation was carried out using the TSCORE evaluation method on a set of seven reference scenes provided by NADC. These scenes were derived from IRAMMP data and were reprocessed to simulate representative backgrounds for a particular airborne IRST sensor. Performance

comparisons were based on a T-90 metric, defined as the peak additive target amplitude required for 90% probability of detection in each reference scene (at a fixed false alarm probability of 3×10^{-5}).

The TSCORE results indicated that the fully-adaptive filter configuration provided an average gain in detection sensitivity of 2.6 dB with respect to the baseline IRST processor. The simplified adaptive filter outperformed the baseline by an average of 1.9 dB, although its T-90 values showed more variation from one reference scene to another. Both adaptive filters appeared capable of providing a more nearly constant (and more predictable) false alarm rate than the baseline processor. In terms of relative computational cost per pixel, the fully-adaptive filter would be over an order of magnitude more complex to implement than the baseline IRST processor, while the simplified adaptive filter would require nearly the same amount of computation as the baseline.

The Phase I effort was successful in demonstrating the application of adaptive spatial filtering to the IRST target detection problem, and in developing a simplified adaptive filter with acceptable detection performance. However, the further development of adaptive algorithms which operate solely in the spatial domain probably cannot be justified by the relatively small performance gains obtained with respect to the existing baseline processor.

2. Results of Phase I Work

2.1 Statement of Problem

The basic problem addressed in the Phase I study was the use of adaptive spatial filtering concepts to enhance the detectability of long-range targets viewed against background clutter by an airborne Infrared Search and Track (IRST) system. For the most part, these stressing threats appear to the IRST sensor as unresolved "point" targets moving against a structured background scene consisting of clouds, blue sky and possibly terrain. The IRST signal processor must be capable of detecting such targets at the lowest possible signal-to-clutter ratios, while strictly controlling the number of false alarms that must be dealt with by higher levels of data processing (e.g., scan-to-scan tracking). The detection signal processing must typically be performed on tens of thousands of pixels in two-dimensional "frames" collected at sensor scan periods on the order of 10 seconds.

Although it is recognized that temporal, spectral and/or polarization discriminants are likely to play an increasing role in the detection processing chains of future systems, near-term IRST systems will continue to rely primarily on spatial processing to perform critical clutter suppression and target detection functions. Thus, the optimization of spatial detection discriminants is a subject of considerable importance. The general areas which must be considered in the design of a practical spatial discrimination scheme are as follows:

- a) the type of filtering to be employed;
- b) filter adaptation to the clutter and noise background;
- c) filter output thresholding (i.e., the detection decision logic); and
- d) performance vs. computation trades.

Of particular interest are relatively *simple* spatial filtering schemes which can provide high detection sensitivity in a wide variety of backgrounds while maintaining a low and nearly-constant false alarm rate.

2.1.1 Existing Approaches

Existing approaches to IRST spatial filtering appear to fall into one of two general categories: 1) fixed filtering schemes and 2) semi-adaptive filter bank techniques. Fixed filtering involves the application of a single linear or nonlinear spatial filter to the entire scan frame. A number of fixed filters which have been considered in the IRST context are discussed in [1]. A single "matched filter" can be designed to provide near-optimum detection performance in a known clutter background, and, for the case of a point target, is often fairly simple to implement. However, the performance of any fixed filter is limited by the inherently nonstationary nature of real-world scenes, where the clutter statistics may change unpredictably from one region to another. Mismatch between the actual background characteristics and those assumed in the design of the filter can result in reduced target sensitivity and/or an increased number of false alarms.

One proposed solution to the mismatch problem is the use of several different filters, each optimized for a particular type of background, in conjunction with a selection logic to decide which filter should be employed in various portions of the scene. A semi-adaptive "filter bank" approach of this type

is described in [2]. The approach is based on partitioning the continuum of background power spectral densities (PSDs) that might be encountered by theIRST sensor into several distinct classes (typically 3 or 4). One class, for example, might represent a nearly white background consisting of blue sky or sensor noise, while another could correspond to a region with extended clouds. For each background class, one designs a suitable matched filter to detect the known target signature. The clutter partitioning and filter design process is carried out entirely off-line, using a training set of representative IR scenes. On-line filter selection is implemented on a pixel-by-pixel basis by means of a sliding-window clutter average. Although ideally the filter choice would be based on the measured PSD (or spatial covariance) of the local background, it has been found empirically that the background type can often be reliably predicted by much simpler surrogate measures, such as local rms intensity.

A precomputed filter bank solution is attractive because it provides a limited degree of adaptivity with a relatively simple computational structure. However, its performance ultimately depends on the number and types of filters selected and on the extent to which the training data used in the filter design are representative of the wide range of clutter conditions encountered in the real world.

2.1.2 Fully-Adaptive Spatial Filtering

A more robust approach to spatial discrimination is motivated by the general theory of statistical hypothesis testing, which has been applied to signal detection problems by Kelly [3] and Reed et. al. [4-6]. In this approach, a likelihood ratio decision statistic is used to test for the presence of a hypothesized target signal at each pixel location. Application of the test over an entire frame constitutes a nonlinear spatial filter whose parameters are continuously adjusted, as a function of the locally measured background clutter covariance, to optimize detection performance for a specified target response.

A fully-adaptive filter offers several potential advantages with respect to the spatial filter bank technique. Most important is the fact that the filter parameters can be locally adjusted to closely match the background statistics actually encountered in each scan frame, rather than those encountered in an off-line training sample. In highly nonhomogeneous clutter, one would expect a fully-adaptive filter to outperform even a well-designed bank of precomputed filters, so long as its clutter estimates are sufficiently accurate and adapt rapidly enough to the nonstationary background.

A second advantage is that the output detection statistic from an adaptive filter can be automatically normalized to provide true constant false alarm rate (CFAR) performance in an unknown Gaussian clutter background. This simplifies the detection logic by eliminating the need for separate threshold estimation procedures, and provides for improved false alarm regulation.

The main disadvantage of a fully-adaptive filter is the computational cost of implementing it. This added complexity results primarily from the need to estimate and invert a clutter covariance matrix at each pixel location in the frame. This problem was clearly recognized in our Phase I proposal, where we noted that applying relatively simple prefilters to IR scenes often served to approximately "diagonalize" (or locally whiten) the background covariance and greatly simplify the adaptive filter implementation. The use of linear and nonlinear prefilters for this purpose was therefore an important topic of investigation under the Phase I study.

2.2 Phase I Technical Objectives

With the above background in mind, the fundamental questions addressed by our Phase I study can be briefly summarized as follows:

- 1) Does fully-adaptive spatial filtering provide a significant improvement in IRST point-target detection performance, compared to a precomputed filter bank approach?
- 2) Can this improved performance be obtained at a reasonable cost in terms of overall computational complexity?

The baseline for comparison was a precomputed spatial filter bank developed by GE and NADC for IRST detection processing.

2.3 Technical Approach

Our technical approach to the Phase I study was to formulate and implement an adaptive spatial filtering approach to IRST point target detection, and to compare its performance and computational costs with those of an existing baseline filter bank. A robust adaptive filter design was first selected based on the theory of statistical hypothesis testing. This filter, which is designed for the detection of a known target response in a Gaussian clutter background, is fully adaptive to measured clutter statistics and can be shown to provide constant false alarm rate (CFAR) performance.

It has been empirically verified that optical and IR scenes can often be characterized by nonstationary Gaussian random processes with rapidly varying means and more slowly varying covariance functions [7]. For this reason, an adaptive spatial filter is generally preceded by a high-pass prefilter that converts the arbitrary background into one which is approximately zero-mean Gaussian with locally stationary second-order statistics. Since the prefilter turns out to have a significant impact on adaptive filter performance and complexity, prefilter selection trades were investigated in some detail.

Although theoretical predictions of adaptive filter detection performance can be obtained for the ideal case of a Gaussian clutter background, it was felt that a more realistic comparative evaluation should be made using a set of seven actual IRST reference scenes provided by NADC. Therefore, the detection sensitivity of both the IRST baseline processor and two candidate adaptive filtering approaches was measured for each reference scene using a variation of the TSCORE analysis method [1]. Estimates of computational complexity were also obtained for each algorithm so that performance/cost trades could be evaluated.

In the following sections, we discuss the baseline and adaptive filtering methods which were considered, the role of spatial prefiltering in adaptive filter implementation, the selection of candidate adaptive filter configurations, and the results of a comparative evaluation of detection performance and computational complexity.

2.4 Baseline Spatial Filter Bank

2.4.1 Baseline Processor Description

A block diagram of an IRST spatial filter bank scheme developed by GE under NADC sponsorship is shown in Figure 1. A 21x7 pixel mean removal filter is first applied to the frame video to subtract the local background dc component. Then, for each pixel location in the frame, one of three different 7x5 linear filter kernels is selected and applied to the prefiltered data. The filter selection is based on the estimated clutter magnitude in a 21x7 pixel window centered at each pixel.

The detection logic consists of comparing the rectified spatial filter output at each pixel against an adaptive threshold computed as a constant threshold factor k times the locally-averaged output clutter magnitude. The local absolute clutter average is computed by a background normalizer consisting of a sliding-window of outer dimensions 21x7 pixels and inner dimensions 7x5 pixels. This allows the effective threshold level to increase in areas of high clutter (such as cloud edges) to control the number of false alarms. The inner window (or "hole") is used to avoid target self-thresholding effects.

This processor, which had been implemented and evaluated in previous studies, was adopted as a baseline for the performance and computation load comparisons made during the Phase I study. Specific filter weights and thresholds were provided to Space Computer Corporation by NADC, and are summarized in Table 1.

2.4.2 Application to IRST Point Target Detection

The application of the baseline filtering scheme to IRST target detection is demonstrated by processing Reference Scene 3 with injected targets. Figure 2(a) shows a grey-scale plot of Scene 3; one of seven such scenes provided by NADC for algorithm evaluation. The frame shown has dimensions of 475x118 pixels and contains a representative background of long-wave IR cloud and sky clutter for a particular airborne IRST system. Gaussian sensor noise (1 LSB rms) was also added to the scene to simulate the noise floor in the system of interest.

The ten synthetic point target responses shown in Figure 2(b) were injected by simple addition into the background scene in Figure 2(a) at a peak amplitude of +50. The target responses were calculated from the following 2-D Gaussian point-spread function:

$$s(x,y) = A \cdot \exp\left\{-\frac{1}{2}\left[\left(\frac{x-\bar{x}}{\sigma_x}\right)^2 + \left(\frac{y-\bar{y}}{\sigma_y}\right)^2\right]\right\} \quad (1)$$

σ_x = horizontal blur radius = 0.926 pixels

σ_y = vertical blur radius = 0.726 pixels

\bar{x} = target position in horizontal frame dimension

\bar{y} = target position in vertical frame dimension

A = additive target amplitude

Figure 3(a) plots the output of the 21x7 mean removal filter for Scene 3 plus targets (the first step in the processing sequence of Figure 1). This prefiltering operation enhances point target responses,

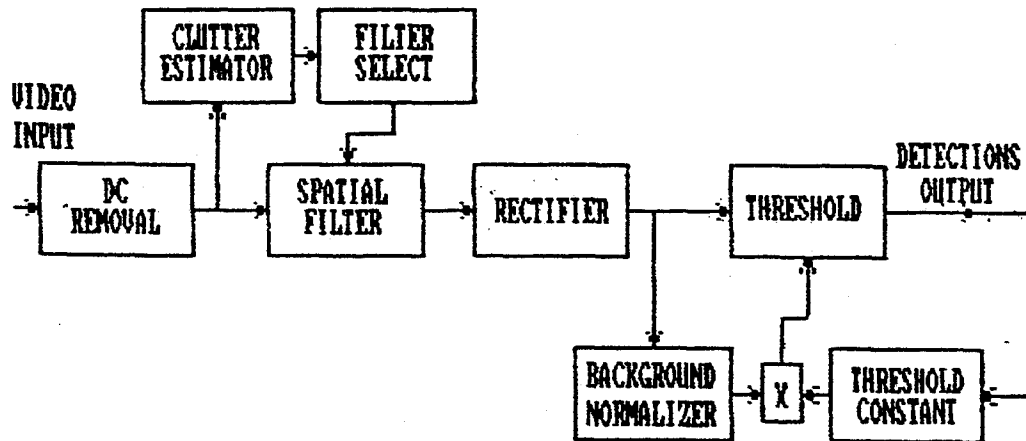


Figure 1. Baseline IRST Detection Processor

Table 1. Baseline Processor Parameters

Pfa specification: 3×10^{-5}
 Pd specification: 90%

Local DC Remove:
 Outer window: 21 az X 7 el pixels
 No inner window

Clutter Estimator:
 Outer window: 21 az X 7 el pixels
 No inner window
 Threshold 1: 6.5
 Threshold 2: 28.0

Spatial Filter

Filter #1 coefficients:

6.41962E-2	1.43601E-2	-3.09115E-2	-0.12691	-2.84954E-2	1.24690E-2	7.60164E-2
-1.65398E-2	-6.97523E-3	-1.30432E-2	8.77677E-2	-4.95103E-3	2.75489E-3	-5.70457E-2
-0.19425	-8.09027E-2	7.60695E-2	0.48568	9.09153E-2	-7.71294E-2	-0.19818
-4.33161E-2	1.07323E-2	7.09682E-3	4.91853E-2	-3.35438E-2	-1.35654E-2	5.10128E-3
8.23992E-2	9.56220E-3	-3.14281E-2	-0.13852	-2.48941E-2	2.05879E-2	6.78911E-2


Filter #2 coefficients:

4.95556E-2	-7.74260E-3	6.17393E-3	-9.66699E-2	7.84275E-3	-1.34985E-2	5.56884E-2
2.60009E-2	1.75524E-2	-7.26968E-2	2.74039E-2	-5.09801E-2	3.30178E-2	1.09579E-2
-0.19790	-8.12331E-2	-1.63643E-2	0.57084	1.14468E-2	-7.72364E-2	-0.20355
1.88873E-2	3.71829E-2	-1.77961E-2	-1.40784E-2	-9.15137E-2	6.70704E-3	5.16049E-2
5.95764E-2	-2.13615E-2	6.05144E-3	-0.10340	2.14771E-2	-1.20005E-3	4.14710E-2


Filter #3 coefficients:

2.05683E-3	-6.03061E-3	8.00534E-2	-4.5592E-2	7.67835E-2	-5.76763E-2	2.33199E-3
4.81053E-2	0.13576	-0.17468	-1.57866E-2	-0.14745	0.10188	6.16529E-2
-0.15051	8.11574E-2	-0.37235	0.83725	-0.32042	7.26099E-2	-0.16139
6.71539E-2	6.83343E-2	-6.13683E-2	-7.87872E-2	-0.17814	5.12027	7.07233E-2
2.72921E-3	-5.94086E-2	7.37562E-2	-5.82136E-2	0.10786	-6.17016E-2	-7.72587E-3

Background Normalizer:
 Outer Window: 21 az X 7 el pixels
 Inner Window: 7 az X 5 el pixels




(a)




(b)

Figure 2. (a) NADC Reference Scene 3.
 (b) Ten Gaussian Point Targets Injected into Scene 3
 at Peak Amplitude +50.



(a)



(b)

Figure 3. (a) Scene 3 after 21x7 Mean Removal Prefiltering.
(b) Filter Selection for Scene 3.
(White=Filter 1; Grey=Filter 2; Black=Filter 3)

high-frequency clutter features (i.e., edges), and sensor pattern noise. Filter selection, the next step, is based on a 21x7 sliding-window magnitude average over the prefiltered scene. The results of the filter selection process applied to Figure 3(a) are graphically illustrated by Figure 3(b). Areas with relatively low rms clutter, plotted as white in Figure 3(b), are assumed to be dominated by blue sky clutter or sensor noise and are processed by an approximation to a "white noise" spatial matched filter (filter number 1). Regions with very high clutter, shown in black, are processed by a specially-designed colored noise filter (filter number 3). An intermediate filter (filter 2) is used in the remaining "grey" areas of moderate clutter.

Figure 4(a) is a plot of the adaptive spatial filter bank output for the prefiltered scene, after normalization by the local clutter estimate from the GE background normalizer. The filter output is plotted in unrectified form for convenience. Note that the point targets now appear to stand out against the suppressed background.

Applying a two-sided threshold k to the scene shown in Figure 4(a) is equivalent to the detection logic shown in Figure 1. The detected output for such a threshold ($k=10$), which is set just high enough to avoid false alarms, is plotted in Figure 4(b). Comparison with Figure 2(b) shows that all ten targets are detected. The results of adaptive filtering are more readily seen in Figure 5, which shows the signal in one of the target channels (number 68) before processing and just prior to thresholding.


2.5 Adaptive Spatial Filtering

The mathematical theory for the adaptive filtering concepts investigated during the Phase I study has been developed in key papers by E.J. Kelly [3] and I.S. Reed et.al. [5,6]. The general problem is that of detecting the presence of a specified signal shape in a background of Gaussian clutter with unknown statistics, using the theory of optimum statistical hypothesis testing. In this approach, a generalized likelihood ratio detector is formulated to test the hypothesis that the signal is present in a "vector" of lexicographically-ordered pixel observations from the frame. A set of independent reference vectors formed from pixels in the vicinity of this test vector is used to estimate the unknown local covariance of the clutter, and formulate an optimum approximation to the matched spatial filter. A test statistic consisting of a normalized matched filter output is then compared with a fixed threshold for detection. This test statistic, which is a nonlinear function of pixel observations, can be shown to have a probability distribution which is completely independent of the level or structure of the background (when no signal is present). A fixed detection threshold therefore provides a constant false alarm rate (CFAR) in an unknown Gaussian clutter background.


2.5.1 Mathematical Formulation

Consider a vector observation \underline{x} consisting of N lexicographically-ordered pixels taken from a 2-D sub-window of a scan frame. We wish to test this observation for the possible presence of a signal of known shape in a combined clutter and noise background. For sub-pixel target detection the signal can be approximated as additive. Then if the target is present we model the observation to be tested as

$$\underline{x} = A\underline{s} + \underline{n} \quad (2)$$



(a)



(b)

Figure 4 (a) Normalized Filter Bank Output for Scene 3.
(b) Threshold Exceedances in Figure 4(a) for $k=10$.

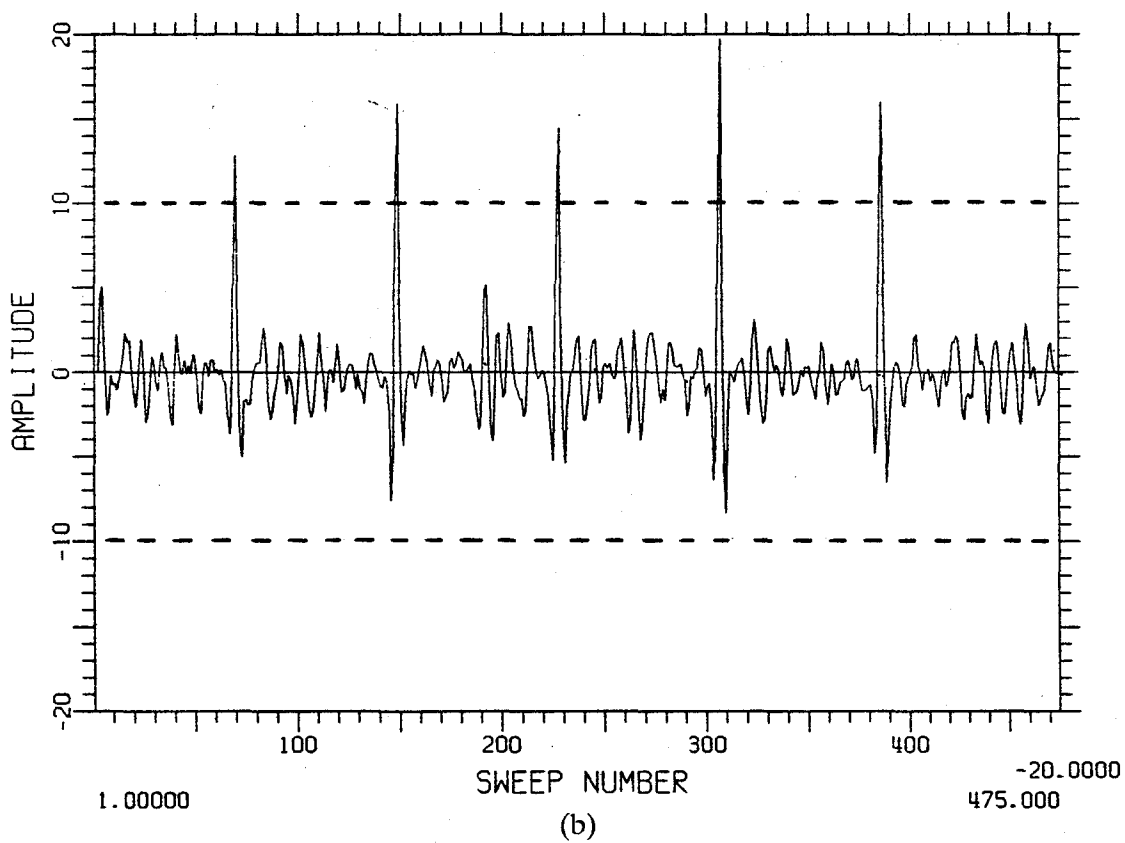



Figure 5 (a) Channel 68 from Scene 3 (with Injected Targets)
 (b) Normalized Filter Output for Channel 68.
 The five threshold crossings correspond to targets.

where \underline{s} is the known target response shape (arbitrarily scaled), A is the target amplitude, and \underline{n} represents the random background. To make the filter derivation tractable, we will assume that the local background (after suitable preprocessing) is jointly Gaussian with zero mean and $N \times N$ spatial covariance matrix

$$\mathbf{M} = E \underline{n} \underline{n}^T \tag{3}$$

where "T" denotes the matrix transpose and "E" indicates an ensemble average. Both the target amplitude A and the background spatial statistics represented by \mathbf{M} are unknown a priori.

Information about the unknown background is available in the form of K N -vectors \underline{x}_k , $k=1, \dots, K$, which constitute "secondary" inputs to the detector. Typically, the \underline{x}_k -vectors are obtained from N -pixel sub-windows in the immediate spatial vicinity of the test window \underline{x} . It is assumed that 1) the target signal, if present, occupies only the pixels in the test window, and 2) the clutter observations from the test window and the surrounding secondary windows are mutually uncorrelated and share the same Gaussian spatial statistics. The degree to which this latter assumption is valid is highly dependent on the frame data, the type of prefilter that has been applied to the data, and the size and arrangement of the various sub-windows. These important issues are discussed further in Section 2.6 below.

In the generalized likelihood ratio detection approach, the respective probability densities of the inputs $(\underline{x}, \underline{x}_1, \dots, \underline{x}_K)$ under each hypothesis are maximized separately over the unknown parameters; in this case, the elements of the $N \times N$ positive definite covariance matrix \mathbf{M} and the target amplitude A . The ratio of these maxima forms a test statistic which can be compared against a threshold for detection. This is a standard technique for statistical hypothesis testing which has been applied to signal detection problems arising in radar, communication and electro-optical systems.

The derivation of the test begins with the joint probability densities of the input observations under the hypotheses H_0 (no signal present) and H_1 (signal present). Based on the assumptions outlined above, the appropriate joint Gaussian probability density under the H_0 hypothesis is

$$P_0(\underline{x}, \underline{x}_1, \dots, \underline{x}_K) = (2\pi)^{-N/2} |\mathbf{M}|^{-1/2} \exp\left\{-\frac{1}{2} \underline{x}^T \mathbf{M}^{-1} \underline{x}\right\} \cdot \prod_{k=1}^K (2\pi)^{-N/2} |\mathbf{M}|^{-1/2} \exp\left\{-\frac{1}{2} \underline{x}_k^T \mathbf{M}^{-1} \underline{x}_k\right\}$$

$$= \left[(2\pi)^{-N/2} |\mathbf{M}|^{-1/2} \exp\left\{-\frac{1}{2} \text{tr}(\mathbf{M}^{-1} \mathbf{T}_0)\right\} \right]^{K+1} \tag{4}$$

where

$$\mathbf{T}_0 = \frac{1}{K+1} \left[\underline{x} \underline{x}^T + \sum_{k=1}^K \underline{x}_k \underline{x}_k^T \right] \tag{5}$$

The latter equality in (4) is obtained by rewriting each inner product of the form $\underline{y} \mathbf{M}^{-1} \underline{y}$ in terms of a matrix trace as

$$\underline{y}M^{-1}\underline{y} = \text{tr}(M^{-1}\underline{y}\underline{y}^T) \tag{6}$$

Under the signal present hypothesis H_1 , the joint probability density also has a Gaussian form similar to that given in (4), except that the test vector \underline{x} now has a nonzero mean and \underline{x} is replaced by $\underline{x} - E\underline{x} = \underline{x} - A\underline{s}$ in accordance with (2). The joint density for H_1 is then

$$p_1(\underline{x}, \underline{x}_1, \dots, \underline{x}_K) = \left[(2\pi)^{-N/2} |\underline{M}|^{-1/2} \exp\left\{-\frac{1}{2}\text{tr}(\underline{M}^{-1}\underline{T}_1)\right\} \right]^{K+1} \tag{7}$$

where

$$\underline{T}_1 = \frac{1}{K+1} \left\{ (\underline{x} - A\underline{s})(\underline{x} - A\underline{s})^T + \sum_{k=1}^K \underline{x}_k \underline{x}_k^T \right\} \tag{8}$$

The generalized likelihood ratio test of H_1 vs. H_0 is formally defined by

$$\lambda(\underline{x}, \underline{x}_1, \dots, \underline{x}_K) = \frac{\max_A \max_M p_1(\underline{x}, \underline{x}_1, \dots, \underline{x}_K)}{\max_M p_0(\underline{x}, \underline{x}_1, \dots, \underline{x}_K)} \underset{H_0}{\overset{H_1}{>}} \lambda_0 \tag{9}$$

where λ_0 is a pre-selected threshold. It is not necessary to maximize the denominator of (9) over A since $p_0(\cdot)$ does not depend on that parameter.

A useful expression for the detection statistic is found by performing the indicated maximizations using the respective probability densities $p_0(\cdot)$ and $p_1(\cdot)$ given in (4) and (7). First, we note that the maximizations over the unknown covariance matrix elements are equivalent to finding the maximum-likelihood estimates of \underline{M} under H_0 and H_1 . Under H_0 , the Gaussian vector observations $(\underline{x}, \underline{x}_1, \dots, \underline{x}_K)$ are independent and identically distributed, and the maximum likelihood estimate of their common covariance \underline{M} is provided by the well-known sample covariance matrix, which in this case is given by \underline{T}_0 in (5). Substituting \underline{T}_0 for \underline{M} in (4) we obtain

$$\max_M p_0(\cdot) = \left[(2\pi e)^{-N/2} |\underline{T}_0|^{-1/2} \right]^{(K+1)} \tag{10}$$

Similarly, for the H_1 case we have

$$\max_M p_1(\cdot) = \left[(2\pi e)^{-N/2} |\underline{T}_1|^{-1/2} \right]^{(K+1)} \tag{11}$$

Using the above results in the likelihood test (9), and raising each side to the $2/(K+1)$ power for convenience, yields

$$\ell(\underline{x}, \underline{x}_1, \dots, \underline{x}_K) = \frac{\max_A |\mathbf{T}_1|^{-1}}{|\mathbf{T}_0|^{-1}} = \frac{|\mathbf{T}_0|}{\min_A |\mathbf{T}_1|} \underset{H_0}{\overset{H_1}{>}} \ell_0 \quad (12)$$

where the new threshold is $\ell_0 = \lambda_0^{2/(K+1)}$.

To perform the remaining minimization over the unknown target amplitude A we first use (5) and (8) to write expressions for \mathbf{T}_0 and \mathbf{T}_1 as

$$(K+1)\mathbf{T}_0 = \underline{x}\underline{x}^T + K\hat{\mathbf{M}} \quad (13)$$

$$(K+1)\mathbf{T}_1 = (\underline{x}-A\underline{s})(\underline{x}-A\underline{s})^T + K\hat{\mathbf{M}} \quad (14)$$

where $\hat{\mathbf{M}}$ is the sample covariance formed from the secondary (signal-free) data vectors only:

$$\hat{\mathbf{M}} = \frac{1}{K} \sum_{k=1}^K \underline{x}_k \underline{x}_k^T \quad (15)$$

The matrix $\hat{\mathbf{M}}$ is nonsingular with probability one if $K > N$ [3] (i.e., if the number of secondary vector outer products being averaged is at least as large as the test vector dimension). Assuming that this condition is satisfied, the determinants of both sides of (13) and (14) can be evaluated to obtain the following expressions for $|\mathbf{T}_0|$ and $|\mathbf{T}_1|$:

$$|\mathbf{T}_0| = \left[\frac{K}{K+1} \right]^N |\hat{\mathbf{M}}| \left[1 + \frac{1}{K} \underline{x}^T \hat{\mathbf{M}}^{-1} \underline{x} \right] \quad (16)$$

$$|\mathbf{T}_1| = \left[\frac{K}{K+1} \right]^N |\hat{\mathbf{M}}| \left[1 + \frac{1}{K} (\underline{x}-A\underline{s})^T \hat{\mathbf{M}}^{-1} (\underline{x}-A\underline{s}) \right] \quad (17)$$

It is straightforward to show that the target amplitude which minimizes (17) is $A = (\underline{s}^T \hat{\mathbf{M}}^{-1} \underline{x}) / (\underline{s}^T \hat{\mathbf{M}}^{-1} \underline{s})$. Substituting this value for A into (17), and using (12), the likelihood test becomes

$$\ell(\underline{x}) = \frac{1}{1 - \frac{1}{K} \frac{(\underline{s}^T \hat{\mathbf{M}}^{-1} \underline{x})^2}{\underline{s}^T \hat{\mathbf{M}}^{-1} \underline{s} \left[1 + \frac{1}{K} \underline{x}^T \hat{\mathbf{M}}^{-1} \underline{x} \right]}} \underset{H_0}{\overset{H_1}{>}} \ell_0 \quad (18)$$

Following the approach of Kelly [3], we define a new test statistic

$$\eta(\underline{x}) = \frac{\frac{1}{K} (\underline{s}^T \hat{\mathbf{M}}^{-1} \underline{x})^2}{\underline{s}^T \hat{\mathbf{M}}^{-1} \underline{s} \left[1 + \frac{1}{K} \underline{x}^T \hat{\mathbf{M}}^{-1} \underline{x} \right]} \quad (19)$$

such that

$$\ell(\mathbf{x}) = \frac{1}{1 - \eta(\mathbf{x})} \tag{20}$$

The generalized likelihood test is then written in the form

$$\eta(\mathbf{x}) \underset{H_0}{\overset{H_1}{>}} \eta_0 \tag{21}$$

where $\eta_0 = 1 - 1/\ell_0$ is a normalized threshold that lies on the interval (0,1).

To obtain a test which is more readily compared with conventional spatial filters, we rewrite (19) in the equivalent form

$$y(\mathbf{x}) = \frac{|\underline{s}^T \hat{M}^{-1} \mathbf{x}|}{\left[\underline{s}^T \hat{M}^{-1} \underline{s} \left(1 + \frac{1}{K} \mathbf{x}^T \hat{M}^{-1} \mathbf{x} \right) \right]^{1/2}} \underset{H_0}{\overset{H_1}{>}} t_0 \tag{22a}$$

$$\hat{M} = \frac{1}{K} \sum_{k=1}^K \mathbf{x}_k \mathbf{x}_k^T \tag{22b}$$

where $t_0 = \sqrt{K} \eta_0$ is the detection threshold. We selected this particular form of the adaptive filter detector for implementation in the Phase I study.

We can interpret $y(\mathbf{x})$ above as a nonlinear spatial filter applied to the test data vector \mathbf{x} , with K secondary vectors $(\mathbf{x}_1, \dots, \mathbf{x}_K)$ used to estimate the covariance matrix \hat{M} and adapt the filter to the statistics of the local background. A target detection is declared if $y(\mathbf{x})$ exceeds a normalized detection threshold t_0 which lies between 0 and \sqrt{K} . The value of this threshold controls the false alarm rate of the test in the absence of target.

2.5.2 Properties of the Adaptive Filter

The mathematical properties of adaptive detectors defined by (22) have been studied for the complex-valued (radar) case by Kelly [3]. With suitable modifications for real-valued data, the same basic properties carry over to electro-optical image processing applications as well. The most important of these properties in terms of adaptive filter performance and implementation are summarized below.

Scale Invariance. The adaptive filter output $y(\hat{\mathbf{x}})$ is insensitive to a scale change in either the signal template \underline{s} or the data vectors $(\mathbf{x}_1, \dots, \mathbf{x}_K)$. This is due to the data and signal dependent normalization factors that appear in the denominator of (22a).

Output Range. The adaptive filter output is a normalized quantity that always lies between zero and \sqrt{K} , where K is the number of secondary vector inputs used to form the covariance estimate \hat{M} . Clearly, the detection threshold t_0 must also be in this range.

It is sometimes more convenient to present the adaptive filter output in its unrectified form (i.e., to omit the $|\cdot|$ operation in the numerator of (22a)). In this case, the filter output range is $\pm\sqrt{K}$ and t_o is interpreted as a two-sided threshold.

CFAR Property. The adaptive filter defined by (22) achieves a constant false alarm rate (CFAR) in a Gaussian clutter background. That is, for a given threshold t_o , the detection test (22a) has a fixed probability of false alarm which is completely independent of the level or structure of the background covariance M . The theoretical relationship for the false alarm probability P_{fa} in terms of the threshold t_o is in fact given by

$$P_{fa} = 1 - A_{K+1-N}(\tau) \tag{23}$$

$$\tau = \left[(K+1-N) \cdot \frac{t_o^2 / K}{1 - t_o^2 / K} \right]^{1/2}$$

where $A_\nu(\tau)$ is the Student-t cumulative distribution function with ν degrees-of-freedom [8, p. 948], N is the dimension of the test vector and K is the number of secondary data vectors. The claimed CFAR property holds since P_{fa} depends only on the threshold and the dimensional parameters N and K .

Figure 6 shows typical curves of P_{fa} vs. the threshold t_o for several combinations of the parameters N and K . The curve labeled $K=\infty$ shown in each plot corresponds to the P_{fa} curve for a single rectified unit-power Gaussian variable. The P_{fa} characteristic for the adaptive filter approaches this standard curve asymptotically as $K \rightarrow \infty$ for any value of N . In this limit, the filter output pdf (under H_o) therefore becomes a rectified standard normal distribution.

Relationship to Matched Filter. The adaptive filter detector reduces to the conventional colored noise matched filter as the clutter covariance estimate becomes "perfect." To see this, let the number of secondaries $K \rightarrow \infty$ so that \hat{M} converges to the true clutter covariance M . Then the detection test shown in (22) can be written as

$$\frac{|\underline{s}^T \hat{M}^{-1} \underline{x}|}{\sigma_o} \underset{H_o}{\overset{H_1}{>}} t_o \tag{24}$$

where σ_o is

$$\sigma_o = \left[\underline{s}^T M^{-1} \underline{s} \right]^{1/2} = \left[\text{rms filter output under } H_o \right] \tag{25}$$

The numerator of the left side of (24) is the classical matched filter for detecting the signal \underline{s} in Gaussian noise of *known* covariance M . Note that this same generic form also appears in the numerator of the adaptive filter output in (22a), with the estimate \hat{M} replacing M .

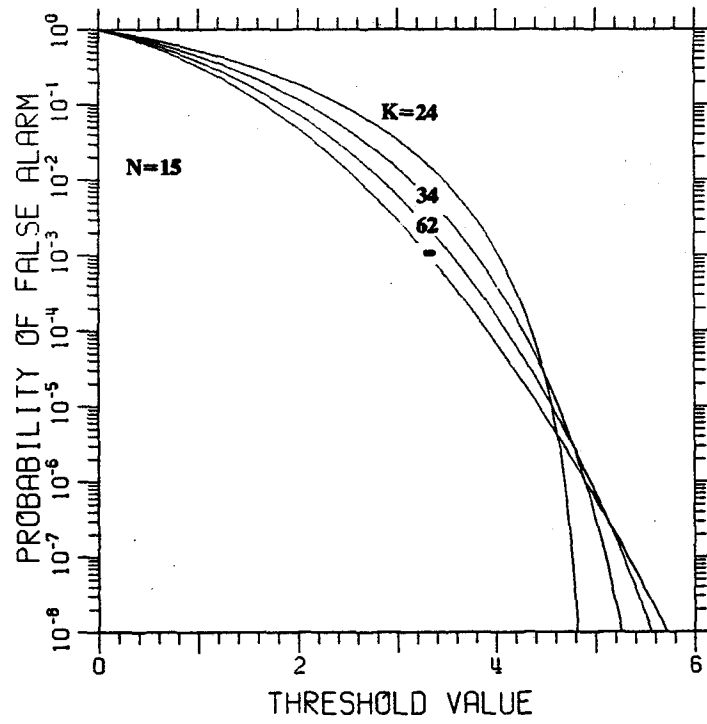



Figure 6. P_{fa} Curves for Several Adaptive Detector Configurations

The classical matched filter output, when properly normalized by its standard deviation σ_o in (25), is compared against a fixed threshold t_o for detection. The role of the denominator of the adaptive filter output (22a) is entirely analogous to that of σ_o : it normalizes the output of the *estimated* matched filter $|\underline{s}^T \hat{\mathbf{M}}^{-1} \underline{x}|$ so that its probability distribution (under H_o) does not depend on either the signal template or the clutter covariance. This CFAR property is an important practical feature that does not hold for most spatial filters. Primarily for this reason, the use of a filter derived from the generalized likelihood principle is conceptually superior to ad-hoc approaches to spatial adaptivity, such as simply plugging a local estimate of \mathbf{M} into the standard matched filter equation.

2.5.3 Application to IRST Point Target Detection

A block diagram of the adaptive filter processor is shown in Figure 7. The input frame is first passed through a high-pass prefilter to convert the arbitrary background to one that is more nearly approximately by zero-mean Gaussian statistics (the role of prefiltering is discussed further in Section 2.6). Adaptive spatial filtering is then implemented in sliding-window fashion on the prefiltered data. At each pixel location, the filter output is computed according to the numerator of (22a), based on the locally-measured clutter covariance and the shape of the desired signal. This output is then normalized and compared against a fixed CFAR threshold for detection.

Figure 8 shows a typical sliding window configuration for an adaptive spatial filter, with a 5x3 pixel ($N=15$) "test window" \underline{x} surrounded by a total of $K=34$ 5x3 "secondary" or "CFAR" windows \underline{x}_k arranged in a 7x5 block pattern (we refer to this as a 5x3/7x5 configuration). The underlying assumptions here are that 1) the target signal, if present, occupies only the pixels in the filter window, and that 2) the clutter observations in the filter window and the surrounding CFAR windows are mutually uncorrelated and share the same Gaussian spatial statistics (after prefiltering).

At each filter position, pixel data from the 34 surrounding CFAR windows are used to form a set of secondary data vectors \underline{x}_k , each of length 15. The outer products of these vectors are averaged to estimate the unknown clutter covariance matrix for the 15 pixels of the filter window. This covariance matrix is then inverted to calculate the filter weights needed to implement the adaptive filter and properly normalize its output for CFAR detection. The estimation and inversion of the covariance matrix for each pixel location turns out to be the major driver of computational complexity for the fully-adaptive filter.

The application of the 5x3/7x5 adaptive filter to IRST point target detection is demonstrated using Reference Scene 3 with ten inserted Gaussian point targets of peak amplitude 50. The scene with the injected targets shown in Figure 1(b) is plotted in Figure 9(a). The adaptive filter is matched to a 5x3 pixel signal vector \underline{s} obtained from the Gaussian point spread function (1).

The first step, spatial prefiltering, is implemented here with a sliding-window mean-removal filter having outer dimensions of 7x5 pixels and inner dimensions (or "hole") of 3x3 pixels. This filter does a reasonable job of removing the local background mean while avoiding self-suppression of the desired target response. The prefilter output frame is plotted in Figure 9(b). Point target responses are now evident but a significant amount of structured clutter remains, particularly in the vicinity of cloud edges.

The unrectified adaptive filter output for this prefiltered frame is plotted in a grey-scale format in Figure 10(a), with black and white corresponding to the boundaries of the filter output range ($\pm\sqrt{34}$). Since

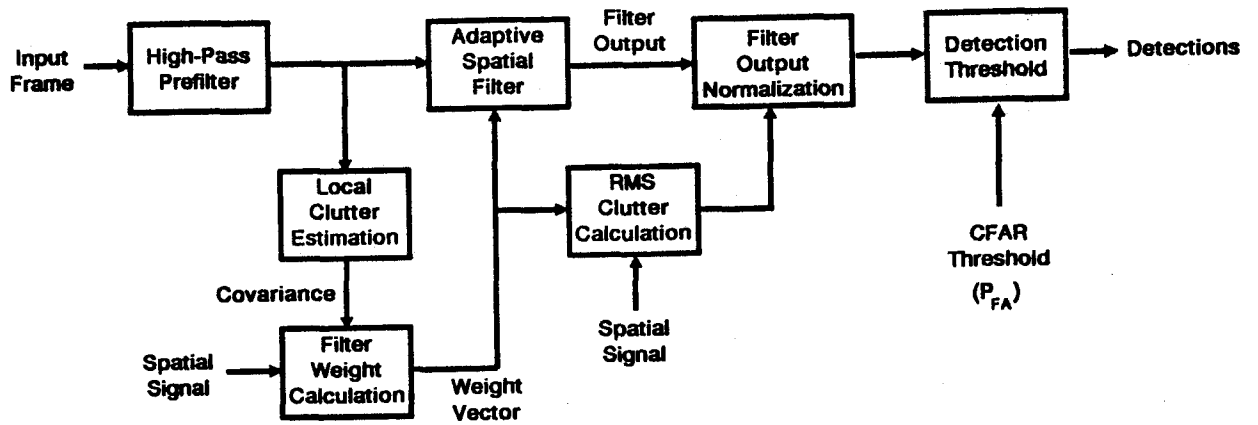



Figure 7. Adaptive Spatial Filter Processor

5 x 3 / 7 x 5 Configuration




**5 x 3 Filter Window for Target Centered at ■
(15 Pixels --> 15 x 15 Covariance Matrix)**




**5 x 3 CFAR Windows Arranged in 7 x 5 Pattern
(34 Windows for Clutter Covariance Estimation)**

Figure 8. Adaptive Filter Window Configuration




(a)




(b)

Figure 9 (a) Reference Scene 3 with 10 Injected Targets of Amplitude +50.
 (b) Scene 3 after 7x5/3x3 Mean Removal Prefiltering



(a)



(b)

Figure 10(a) Normalized Adaptive Filter Output for Scene 3.
(b) Threshold Exceedances in Figure 10(a) for $t_0 = 5$.

this represents the output of a complex nonlinear adaptive filter, it is somewhat misleading to use the relative intensities in this "scene" to assess the detectability of targets prior to thresholding. Nonetheless, the targets are readily detected from this output. Figure 10(b) shows the locations of threshold exceedances in the frame at a theoretical P_{fa} of 10^{-6} ($t_o = 5.0$); these correspond to the 10 injected targets shown in Figure 2(b).

The results of adaptive filtering are more easily viewed in Figure 11, where plots of signal magnitudes in target channel 68 are shown before and after adaptive filter processing. The solid lines in Figure 11(b) represent the filter output bounds of $\pm\sqrt{34}$. The fixed detection threshold is shown by the dotted lines.

2.5.4 Theoretical Performance Prediction

The detection performance of the adaptive filter in a Gaussian clutter environment can be explicitly calculated by specializing the methods developed in [3] to the case of real-valued data. Only four parameters affect this performance:

N = the dimension of the test and secondary vectors (the number of pixels in the test window);

K = the number of secondary vectors (windows) used to estimate the clutter covariance matrix;

P_{fa} = the design probability of false alarm

SCR = a local signal-to-clutter ratio = $\{E_x\}^T M^{-1} \{E_x\} = A^2 \mathbf{s}^T M^{-1} \mathbf{s}$

Note that the generalized SCR defined here reduces to the more familiar point-target SCR expression A^2/σ^2 if the background is white (i.e., $M = \sigma^2 \mathbf{I}$) and the signal shape template is normalized to unit power ($\mathbf{s}^T \mathbf{s} = 1$).

Sample curves of probability of detection (P_d) vs. SCR are shown in Figure 12 for $N=15$ (e.g., a 5×3 pixel test window) and several different values of K . The P_{fa} is fixed at 10^{-6} . The curve labeled $K=\infty$ represents the performance of a perfectly-matched spatial filter where the clutter statistics are completely known in advance. The SCR differences between this curve and the various adaptive filter curves indicate the magnitude of the total CFAR losses incurred in using finite samples of data to estimate the clutter covariance matrix and normalize the filter output for thresholding. These losses can be quite large compared with typical CFAR losses for scalar detectors (where $N=1$).

Calculated CFAR losses at $P_d=0.9$ and $P_{fa}=10^{-6}$ for some $N=15$ and several values of K are provided in Table 2 below. Clearly, this source of loss can be reduced to an arbitrarily low level if a sufficient number K of secondary windows (having the same statistics as the test window) can be provided to the detector. In practice, however, the use of a very large value of K (many secondary windows) results in a spatial filter that is relatively slow to adapt to clutter variations in a nonhomogeneous scene. The CFAR loss discussed here is only one of several issues to be considered in selecting a window configuration for adaptive spatial filtering.


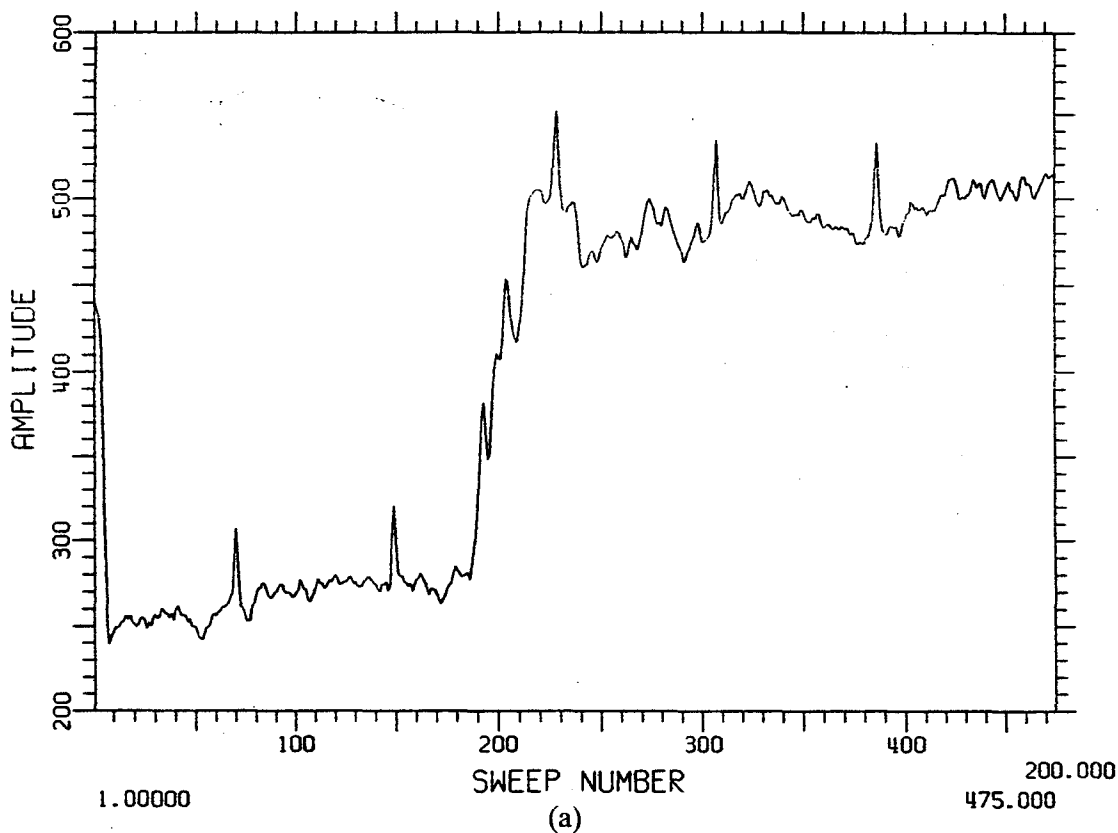


Figure 11 (a) Channel 68 from Scene 3 (with Injected Targets)
(b) Normalized Adaptive Filter Output for Channel 68.
The five threshold crossings correspond to targets.




Figure 12. Probability of Detection vs. SCR for Several Adaptive Filter Configurations ($P_{fa} = 10^{-6}$)

The entries in Table 2, along with additional calculations shown in [3], indicate that the CFAR penalty becomes very large if K drops below about $2N$. As a rule of thumb, the number of secondary windows must be at least $2N$ to hold the total CFAR loss to around 6 dB or less. This is a large penalty to pay for the benefits of fully-adaptive operation.

Table 2. Adaptive Filter CFAR Loss

Configuration	N	K	CFAR Loss
5x3/5x5	15	24	16.2 dB
5x3/7x5	15	34	6.4
5x3/9x7	15	62	2.8
5x3/ ∞	15	∞	0.

2.6 Spatial Prefiltering

The spatial prefiltering function is a critical part of the adaptive filter processing chain. The primary function of the prefilter, which is generally chosen to have a high-pass characteristic, is to remove the local mean level of the background and "Gaussianize" the residual clutter component. This is often done with simple sliding-window mean removal operators, but it can also be approximated by nonlinear filtering (such as sliding-window median removal).

It has been empirically verified that many natural IR scenes can be characterized as Gaussian random processes with a rapidly varying mean and a more slowly varying covariance [7]. To the extent that this is true, a high-pass prefilter will tend to convert an arbitrary nonhomogeneous background into one which is approximately zero-mean Gaussian with locally-stationary second-order statistics. Such a background is consistent with the assumptions under which the adaptive filter was designed. High-pass filtering also has an intuitive appeal since it enhances the responses of point targets embedded in a background of extended features (e.g., clouds).

Prefiltering has a strong affect on the spatial covariance of the frame data which will be processed by the adaptive filter (i.e., the degree to which neighboring pixels are correlated with one another after prefiltering). Generally speaking, the smaller the prefilter window, the more the background tends to decorrelate. Background decorrelation is desirable for the following reasons:

- 1) It permits the use of smaller adaptive filter windows, providing faster adaptivity in nonstationary clutter and reducing the real-time computation load; and
- 2) It may allow much simpler adaptive filters, based on a diagonal approximation to the clutter covariance matrix, to be successfully employed for target detection.

The prefilter also affects the magnitude and shape of additive target signals in the frame. Large high-pass prefilter kernels generally cause less point target attenuation than smaller ones. Selecting a good prefilter therefore involves balancing the desire for robust local mean suppression and maximal

background decorrelation against the need to avoid excessive target attenuation. These two goals are conflicting and necessarily imply a tradeoff.

Two generic types of prefilters were examined during the Phase I study: 1) linear high-pass filters based on sliding-window mean removal and 2) nonlinear filters based on sliding-window median removal. Their properties are discussed in turn below.

2.6.1 Linear High-Pass Filters

Linear sliding-window mean removal operators are used extensively in image processing. They are very simple to implement: each pixel in the frame is replaced by its original value minus the local average computed in a surrounding window.

The key parameters examined in this study were the window dimensions, which determine the spatial frequency passband of the filter. We considered estimation windows of relatively small spatial extent, such as 3x3, 5x3 and 7x5 pixels (with and without inner windows or "holes"). Such filters do a reasonable job of tracking the rapidly-varying local mean of the a typical IR background. Given the high degree of background variation in the NADC Reference Scenes, the 21x7 window used in the baseline filter design was deemed to be excessively large.

The general performance characteristics of mean removal filters with respect to background Gaussianization, background decorrelation, and target attenuation are discussed below.

Background Gaussianization. A key function of the prefilter is to convert the arbitrary scene background into one that is reasonably well characterized by zero-mean, locally Gaussian clutter statistics. This allows the considerable body of detection theory that has been developed for Gaussian processes to be successfully applied to the IRST detection problem.

A simple criterion for selecting a linear prefilter which is "best" in this sense has been investigated by Reed et.al. [4]. Their approach is to select the window size that minimizes the normalized third moment of the output clutter histogram, which is zero for the Gaussian distribution. Various tests for Gaussian statistics, such as chi-square goodness-of-fit or the Kolmogorov-Smirnov test on the sample cdf, can also be employed.

A related and perhaps more important criterion is the degree to which a prefilter suppresses the structure in the scene, particularly along edges corresponding to clutter boundaries. Visible edges in a prefiltered scene indicate the presence of deterministic fine structure that cannot be characterized by a zero-mean random clutter process. These edges are a major source of false alarms in IR surveillance sensors.

For the NADC Reference Scenes, we found that smaller linear filter windows were consistently superior to larger ones in terms of the above criteria. The qualitative effect of prefilter size is illustrated by the plots shown in Figures 13, 14 and 15. Figure 13 shows Reference Scene 3 along with its amplitude histogram. Figures 14 and 15 show the prefilter outputs obtained from Scene 3 using 7x5/3x3 and 3x3 mean removal windows, respectively (these frames are plotted on the same grey-scale for comparison). Although structured clutter is present in both cases, the smaller 3x3 window provides better overall suppression and produces considerably less "smearing" of the cloud edges. Histograms for the two output frames are also shown. Both prefilters convert the original background histogram in Figure 13(b) to a

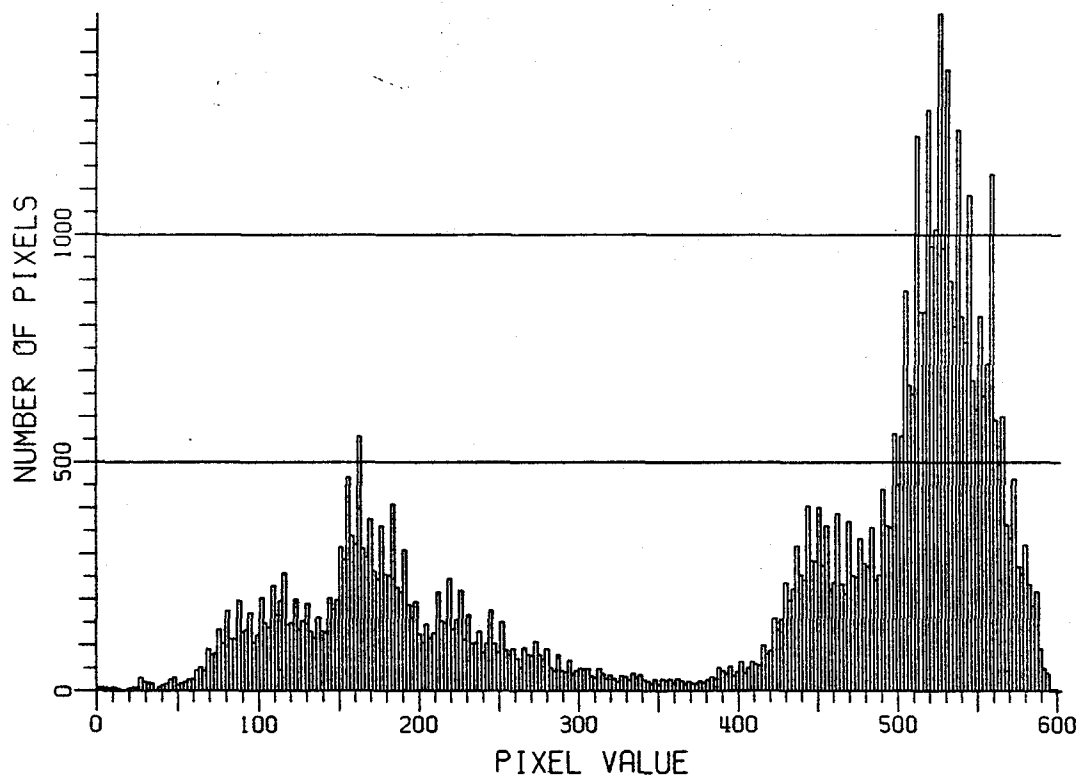



Figure 13. Reference Scene 3 and its Histogram.

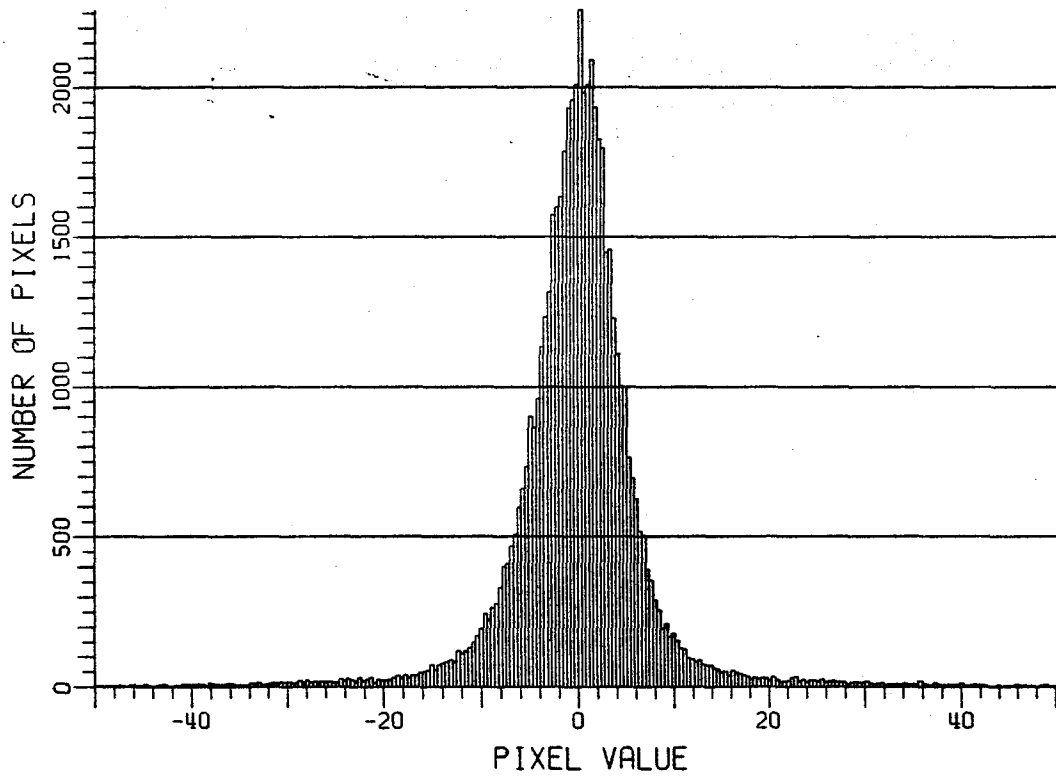



Figure 14. Scene 3 and Histogram after $7 \times 5 / 3 \times 3$ Mean Removal Prefiltering.





Figure 15. Scene 3 and Histogram after 3x3 Mean Removal Prefiltering.

unimodal, symmetric zero-mean distribution. However, the output histogram for the 7x5/3x3 filter in Figure 14(b) exhibits relatively long tails due to the bright edges. The 3x3 output histogram in Figure 15(b) is closer to a true zero-mean Gaussian distribution in this sense. Moreover, the rms clutter level from the 3x3 prefilter is considerably lower due to that filter's larger stopband about dc spatial frequency.

Clutter Decorrelation. The spatial correlation properties of the prefilter output directly affect the design of the adaptive filter, since the spatial extent of the adaptive filter test window must generally be chosen to match the actual correlation length of the clutter. This ensures that pixels with significant mutual correlation are processed together in the same 2-D matched filter window, and that the clutter observations in adjacent "secondary" windows, which are used for clutter averaging, are approximately uncorrelated with one another. Maximal decorrelation of the background is desirable since it permits the use of small adaptive filter windows, which require less computation and adapt faster in nonhomogeneous clutter. The best possible result is total decorrelation of the background. In this case, the clutter covariance is diagonal and the matrix inversion required for adaptive filter implementation is greatly simplified.

Not surprisingly, we found that the clutter increasingly decorrelates as the prefilter window gets smaller. The general trend is illustrated by processing Scene 3 with mean-removal filters of dimensions 21x7 (the GE baseline), 7x5/3x3, and 3x3. Figures 16, 17 and 18 show contour plots of the normalized 2-D autocorrelation functions for the output scenes from each of these prefilters. Approximate 1/e correlation lengths and minimum adaptive filter window dimensions are summarized in Table 3 below.

Table 3. Clutter Correlation After Linear Prefiltering
(Reference Scene 3)


Prefilter Dimensions	Clutter Correlation Length (Pixels)		Minimum Adaptive Filter Window
	Horizontal Dimension	Vertical Dimension	
21x7	4.0	2.5	9x5
7x5/3x3	2.7	1.6	5x3
3x3	2.8	1.0	5x3

The decorrelation obtained with the large 21x7 window is clearly inferior to that of the two smaller windows, which appear to be approximately equivalent in this respect. The correlation lengths obtained with the smaller 7x5/3x3 and 3x3 prefilters suggest that an adaptive filter window on the order of 5x3 pixels could be used. However, neither prefilter decorrelates the background to the point where adjacent pixel correlations can be completely ignored. This apparent limitation of linear prefiltering prompted us to examine nonlinear filtering as an alternative method of background preprocessing (see below).

Target Attenuation. A potential drawback of mean-removal prefilters is the unwanted attenuation of desired point target responses. Smaller prefilter windows tend to provide better clutter suppression, but

CONTOURS FROM 0.00 TO 1.00 BY 0.10

BL1B03NPPCOR.SCC




BL1B03N.SCC + QUANTIZED UNIT POWER NOISE
LOCAL MEAN REMOVED: 21 X 7 / 0 X 0 WINDOW
BL1B03NN.SCC
17 BY 17 WINDOW STARTS AT SAMPLE 56, RECORD 56

Figure 16. Normalized Autocorrelation of Scene 3 after 21x7 Mean Removal Prefiltering.

CONTOURS FROM 0.00 TO 1.00 BY 0.10

BL1803NNMCCOR.SCC

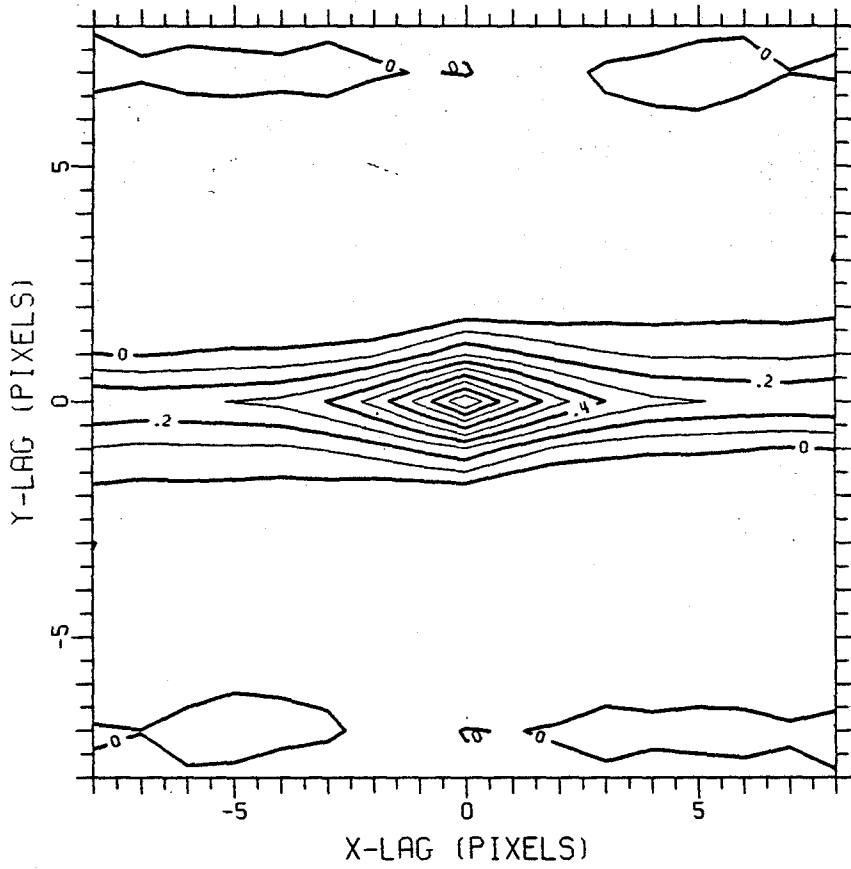


BL1803N.SCC + QUANTIZED UNIT POWER NOISE
LOCAL MEAN REMOVED: 7 X 5 / 3 X 3 WINDOW
BL1803NN.SCC
17 BY 17 WINDOW STARTS AT SAMPLE 56, RECORD 56

Figure 17. Normalized Autocorrelation of Scene 3 after 7x5/3x3 Mean Removal Prefiltering.

CONTOURS FROM 0.00 TO 1.00 BY 0.10

BL1B03NNM3COR.SCC



BL1B03N.SCC + QUANTIZED UNIT POWER NOISE
LOCAL MEAN REMOVED: 3 X 3 / 0 X 0 WINDOW
BL1B03NN.SCC
17 BY 17 WINDOW STARTS AT SAMPLE 56, RECORD 56

Figure 18. Normalized Autocorrelation of Scene 3 after 3x3 Mean Removal Prefiltering.

do so at the expense of additional point target suppression. Table 4 shows the peak amplitude attenuation incurred with various linear prefilters for the desired Gaussian point target response given in (1):

Table 4. Point Target Amplitude Attenuation for Linear Prefilters

Prefilter Dimensions	Peak Amplitude Attenuation
21x7	0.97
7x5/3x3	0.97
7x5	0.88
5x3	0.73
3x3	0.58

Although very large windows cause little target attenuation, they perform so poorly in other respects as to be of little interest. Very small prefilters (e.g., 3x3) are good for suppressing and decorrelating structured backgrounds, but cause significant target self-suppression. Target attenuation is best mitigated through the use of a moderate-sized outer window with a small inner window that excludes the pixels near the window center from the local mean estimate. For example, the 7x5/3x3 linear prefilter discussed in several previous examples represents a fairly good compromise between the conflicting requirements for robust clutter suppression and minimal target attenuation.

2.6.2 Nonlinear High-Pass Filters

Simple nonlinear high-pass filters based on sliding-window median removal operators were also investigated under Phase I. These filters work in basically the same way as the linear prefilters discussed above, except that the sample median in a small window (rather than the mean) is estimated and subtracted from the center pixel.

Median filters are well-known in the image processing community for their ability to preserve sharp edges and other high-frequency artifacts of a scene. The motivation for using local median subtraction for IRST preprocessing was the potential for better suppression and decorrelation of structured clutter. A key problem with mean-removal filters is the edge smearing caused by averaging across clutter boundaries. Our hope was that the improved edge estimates provided by local median filters would allow such edges to be suppressed more effectively by subtraction.

Although we examined various window sizes for median filtering, only the very small windows were found to offer an interesting alternative to mean-removal. A particularly effective window has minimal outer dimensions of 3x3 with a single-pixel "hole" in the center (a 3x3/1x1 configuration). The hole is useful for two reasons. First, it creates a median estimation window with an even number of samples (8). Since the median in this case is defined as the average of the 4th and 5th order statistics, it is never exactly

equal to one of the 8 sample values. This prevents an abnormally large proportion of zero-valued pixels from appearing in the output frame. The hole also reduces attenuation of point targets somewhat.

Background Gaussianization. Since the median is a surrogate estimate of the mean for a locally-stationary clutter background, a median subtraction filter might be expected to have scene normalization properties similar to the mean-removal filters. The result of applying the 3x3/1x1 median removal filter to Scene 3 in Figure 13(a) is shown in Figure 19(a). The histogram of this output frame is shown in Figure 19(b). Although the output pixel values reveal the effects of pixel quantization in the original scene, the envelope of this histogram is certainly very comparable to a zero-mean Gaussian distribution.

The most interesting aspect of median removal filtering is its ability to break up clutter edges. This can be seen by comparing Figure 19(a) with the output of the 3x3 mean removal filter shown in Figure 15(a), which are both plotted on the same grey-scale. Although the edges are still discernable in Figure 19(a), they appear to exhibit much less local structure than before.

Background Decorrelation. The normalized autocorrelation of the filtered frame in Figure 19(a) is plotted in Figure 20. A comparison of this function with the function shown in Figure 18 demonstrates that local median subtraction produces a much sharper correlation peak than a comparable mean subtraction operation. The average correlation of adjacent pixels in Figure 20 is approximately 0.4 and 0.2 in the horizontal and vertical dimensions, respectively. Thus, for all practical purposes, the residual clutter in Scene 3 is nearly decorrelated by the median removal processing. This is largely attributed to the reduced edge structure in the output frame in Figure 19(a). Similar results were observed for several other scenes in the NADC data set.

Target Attenuation. Small-kernel median removal filters obviously suppress point target responses. However, due to the nonlinearity of the median estimate, is not so straightforward to calculate this attenuation. In general, it will depend on the specific background observed in the immediate vicinity of the target, and on the absolute level of the additive target signal itself.

We used a scene-specific simulation approach to characterize the target suppression for the high-pass median filter. For a given reference background (e.g., Scene 3), we injected several hundred point target responses of peak amplitude A_{in} at locations throughout the scene. The target responses were calculated from the Gaussian point-spread function (1), and spaced far enough apart in the frame to avoid multi-target suppression effects. The frame with injected targets was then passed through the 3x3/1x1 median removal prefilter. The peak amplitude for each target in the output frame was measured, and a mean output level A_{out} was calculated from the average of these amplitudes. Repeating the experiment for various A_{in} produced a scene-specific relationship for the average output amplitude A_{out} vs. A_{in} .

An example of such a curve for Scene 3 and a 3x3/1x1 median removal prefilter is shown in Figure 21. The symbols show the locations of actual measurements of A_{out} vs. A_{in} . The solid line is a 2nd-order least-squares polynomial fit to these measurements. For comparison, an equivalent curve for a 3x3/1x1 mean-removal filter is shown by the dashed straight line. Note that the median prefilter causes less target attenuation than the linear prefilter at higher target amplitudes. This occurs because the






Figure 19. Scene 3 and Histogram after 3x3/1x1 Median Removal Prefiltering.

CONTOURS FROM 0.00 TO 1.00 BY 0.10

BL1B03NNM03HCOR.SCC



BL1B03N.SCC + QUANTIZED UNIT POWER NOISE
LOCAL MEDIAN REMOVED: 3 X 3 / 1 X 1 WINDOW
BL1B03NN.SCC
17 BY 17 WINDOW STARTS AT SAMPLE 56, RECORD 56

Figure 20. Normalized Autocorrelation of Scene 3 after 3x3/1x1 Median Removal Prefiltering.

2ND-ORDER LEAST-SQUARES FIT FOR AIN VS. AOUT:
 EXPANSION ABOUT THE POINT AOUT = 2.9185
 B0 = 5.5070
 B1 = 1.5337
 B2 = -0.0014




Figure 21. Average Target Amplitude Attenuation of 3x3/1x1 Median Removal Prefilter in Scene 3.

sample median is less influenced by the presence of a few outliers (e.g., a bright point target in the window) than the sample mean.

2.7 Adaptive Filter Selection

Based on the results of prefilter trade studies and a limited amount of experimentation, we selected two adaptive filter configurations for further evaluation. These two filtering schemes utilized different prefilters and adaptive filter implementations.

Adaptive Filter 1. The characteristics of the first adaptive filter processor are summarized as follows:

Adaptive Filter 1:

- 7x5/3x3 local mean removal prefilter
- 5x3 adaptive filter window (N=15 pixels)
- 5x3 spatial signal matched to Gaussian point-spread
- 7x5 configuration of 5x3 secondary windows (K=34 windows)
- Full covariance matrix processing

The 7x5/3x3 configuration was selected as a baseline linear prefilter because it provided acceptable clutter decorrelation with practically no attenuation of the Gaussian target response. A 5x3 (N=15 pixel) adaptive filter window is the minimum window size that is consistent with the clutter correlation lengths produced by this prefilter on several NADC Reference Scenes. For the test window dimension of N=15, we selected a 7x5 configuration of secondary windows to provide K=34 vector samples for the local covariance estimates (this particular window configuration is shown in Figure 8). From Table 2, the theoretical CFAR loss for this processor in Gaussian clutter is 6.4 dB. A limited set of experiments on Scene 3 indicated that the potential reductions in CFAR loss obtained through the use of additional secondary windows were more than offset by the penalties associated with larger (and more slowly-adapting) filter configurations.

Adaptive Filter 1 employs full 15x15 covariance matrix processing since the pixels in its 5x3 test window are often significantly correlated after mean-removal prefiltering.

Adaptive Filter 2. Adaptive Filter 2 is a simplified adaptive configuration designed to exploit the spatial decorrelation provided by a high-pass median prefilter:

Adaptive Filter 2:

- 3x3/1x1 local median removal prefilter
- 3x3 adaptive filter window (N=9 pixels)
- 3x3 spatial signal matched to Gaussian point-spread
- 7x3 configuration of 3x3 secondary windows (K=20 windows)
- Diagonal approximation to clutter covariance matrix

As noted above, the $3 \times 3 / 1 \times 1$ median high-pass filter decorrelates the residual clutter to the point where its spatial covariance can be approximated by a diagonal (but not necessarily white) matrix. This simplifies the covariance processing for an N -pixel adaptive filter by reducing the number of unique matrix elements to be estimated from $N^2(N^2+1)/2$ to N , and by eliminating a time-consuming $N \times N$ matrix inversion. (However, it does cause scene-dependent target attenuation as discussed in Section 2.6.2.)

The 3×3 adaptive filter window chosen for Adaptive Filter 2 is the minimum window needed to integrate the significant portion of the Gaussian point target response (1). The (diagonal) covariance for this $N=9$ pixel window is estimated from $K=20$ secondary windows arranged in a 7×3 block configuration.

2.8 Algorithm Performance Evaluation

The detection performance of the baseline spatial filter bank and the two candidate adaptive filter configurations was evaluated using a set of seven IR reference scenes provided by NADC. These scenes were derived from data collected under the IRAMMP program and have been reprocessed to simulate the clutter that would be seen by a typicalIRST system. A typical scene, plotted in grey-scale format in Figure 2(a), contains cloud and sky clutter plus additive Gaussian sensor noise.

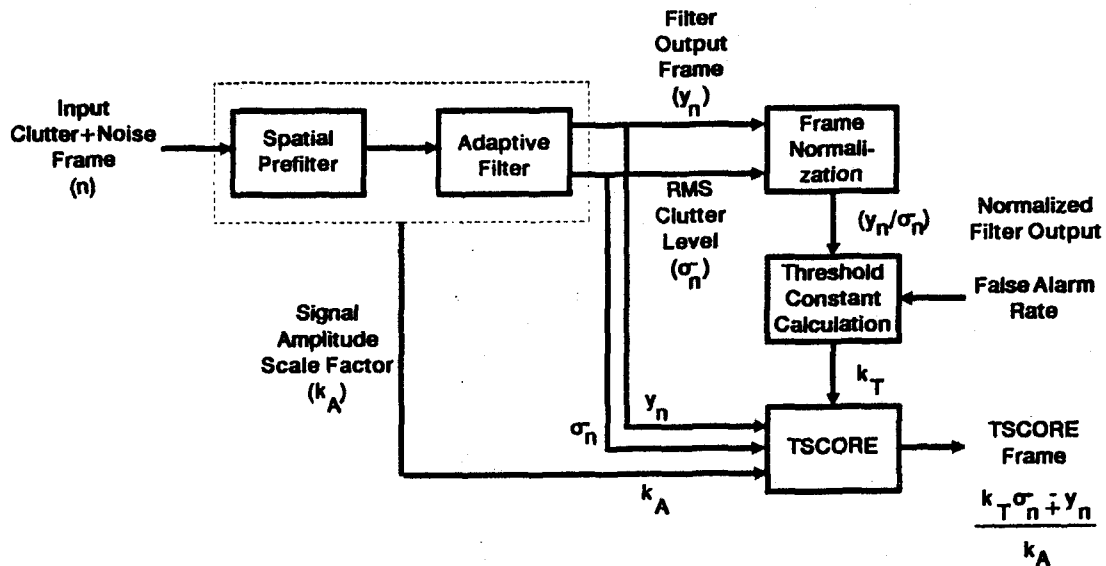
2.8.1 TSCORE Evaluation Methodology

Detection performance was evaluated using the TSCORE method developed by NADC [1]. TSCORE is a general analytical technique for measuring the detection sensitivity of a spatial filter applied to a particular background scene. For every pixel in the prefiltered scene, one calculates the peak additive target amplitude that would be needed to obtain a detection (at a specified average false alarm rate). This calculation must properly account for any scaling that is applied to the target signal by either the prefilter or the adaptive filter. It also implicitly assumes that the threshold level at the target location is independent of the target amplitude.

The output of the TSCORE process is another frame which contains the required target amplitude vs. pixel position. These amplitudes are accumulated in a histogram and integrated to form a cumulative distribution function which gives the probability of detection vs. target amplitude, averaged over a particular scene. The results of a TSCORE analysis are often summarized by the so-called "T-90" figure-of-merit, which is the target amplitude needed to obtain an average detection probability of 90% at a specified false alarm rate.

A block diagram of the TSCORE procedure for a generic adaptive filter is shown in Figure 22. Specific TSCORE calculation methods for both the baseline filter bank and the adaptive filter are developed below.

Baseline Filter Bank. The TSCORE procedure for the GE filter bank technique described in Section 2.4 can be defined by considering what happens to a point target centered at a single pixel in a prefiltered frame. To simplify matters somewhat, assume that the 21×7 mean removal prefilter has a negligible effect



Additive Point Target of Amplitude $\pm A$ is Detected if ...
 Output Signal + Output Clutter $\geq \pm$ Adaptive Threshold

$$k_A(\pm A) + y_n \geq \pm k_T \sigma_n$$

$$A \geq \frac{k_T \sigma_n + y_n}{k_A} = \pm \text{TSCORE}$$

Figure 22. TSCORE Evaluation of a Generic Adaptive Filter

on the amplitude of an isolated point target (practically true), and that the filter selection is relatively unaffected by the presence of a small target in the clutter averaging window. Now define the following 35-element vectors representing 7x5 windows centered at the pixel of interest:

- \mathbf{n} = the clutter background in the window (after prefiltering)
- \mathbf{x} = the target plus clutter observation (after prefiltering)
- \mathbf{s} = the 2-D Gaussian point spread function (normalized to have unit peak magnitude)
- \mathbf{s}_f = the spatial filter weights selected for this pixel

If a target with additive peak amplitude A is present, the observation in the window is given by $\mathbf{x} = A\mathbf{s} + \mathbf{n}$. The output of a linear spatial filter \mathbf{s}_f for the pixel of interest is then

$$y = \mathbf{s}_f^T (A\mathbf{s} + \mathbf{n}) = A\mathbf{s}_f^T \mathbf{s} + \mathbf{s}_f^T \mathbf{n} \quad (26)$$

The GE detection procedure compares the magnitude of y with an adaptive threshold computed by multiplying a local clutter level estimate $\hat{\sigma}$ times a fixed threshold constant k . The clutter estimate is not influenced by the target since it is obtained from an average of pixels outside the 7x5 window. From (26), a detection will be declared at the pixel of interest if the following condition is satisfied:

$$|A\mathbf{s}_f^T \mathbf{s} + \mathbf{s}_f^T \mathbf{n}| > k\hat{\sigma} \quad (27)$$

The TSCORE calculation at a given pixel amounts to solving (27) for the minimum value of A needed for detection. The solution for positive contrast targets ($A > 0$) is given by

$$A_{\min} = \frac{k\hat{\sigma} - F}{S} \quad (28)$$

where

$$F = \mathbf{s}_f^T \mathbf{n} = \text{the spatial filter output for clutter alone}$$

$$S = \mathbf{s}_f^T \mathbf{s} = \text{the signal scaling of the spatial filter}$$

Note that F , S and $\hat{\sigma}$ can be obtained for all pixels in a prefiltered background scene by running that scene through the GE adaptive filter processor. The output is three new "scenes" which are combined according to (28) to produce a TSCORE frame consisting of the minimum detectable target amplitude vs. pixel position at a specified threshold level k .

One final detail is the determination of the threshold constant k . When targets are not present ($A=0$), a false detection will be declared if

$$|\mathbf{s}_f^T \mathbf{n}| = |F| > k\hat{\sigma} \quad (29)$$

at any pixel. To obtain no false alarms, k is chosen such that

$$k \geq \frac{|F|}{\sigma} \tag{30}$$

for every pixel in the frame. More generally, we compute a normalized frame defined by $|F|/\hat{\sigma}$, examine its histogram, and select k to allow no more than a specified number of false detections. The P_{fa} is then approximated by the number of allowed detections divided by the total number of pixels processed.

Adaptive Filter. TSCORE calculation for the adaptive spatial filter defined by (22) is a bit more involved but conceptually quite similar to the procedure discussed above. Even though the filter output is a highly nonlinear function of the observations, the TSCORE method is still valid because the detection threshold setting does not depend on target amplitude. A slight complication is introduced by the need to account for target attenuation in the prefilter, but this can be handled in a straightforward manner.

As before, we define the following vectors having the same dimension as the adaptive filter test window:

- \mathbf{n} = the clutter background in the window (after prefiltering)
- \mathbf{x} = the target plus clutter observation (after prefiltering)
- \mathbf{s} = the 2-D Gaussian point spread function (normalized to have unit peak magnitude)

If the adaptive filter is spatially matched to the target point spread function \mathbf{s} , then its scalar output at a given pixel is defined by the left-hand side of equation (22a):

$$y = \frac{|\mathbf{s}^T \hat{\mathbf{M}}^{-1} \mathbf{x}|}{[\mathbf{s}^T \hat{\mathbf{M}}^{-1} \mathbf{s} (1 + \frac{1}{K} \mathbf{x}^T \hat{\mathbf{M}}^{-1} \mathbf{x})]^{1/2}} \tag{31}$$

Recall that the covariance estimate $\hat{\mathbf{M}}$ is obtained from data in a set of K secondary windows surrounding the filter window, and is therefore unaffected by the presence of a target in the latter window. A detection is declared at the center pixel of the filter window if

$$y > t_0 \tag{32}$$

where t_0 is a fixed CFAR threshold that is set off-line.

When a target of amplitude A is present, the filter window contains the signal $\mathbf{x} = A\mathbf{s} + \mathbf{n}$. Substituting for \mathbf{x} in (31) and using (32) results in a quadratic TSCORE inequality

$$(\gamma_0 S^2)A^2 + (2\gamma_0 SF)A + (F^2 - t_0^2 \sigma_0^2) \geq 0 \tag{33}$$

where $\gamma_0 = 1 - t_0^2/K$ and where

$$F = \mathbf{s}^T \hat{\mathbf{M}}^{-1} \mathbf{n} = \text{the adaptive filter output for the clutter alone}$$

$S = \hat{\mathbf{s}}^T \hat{\mathbf{M}}^{-1} \hat{\mathbf{s}}$ = the adaptive filter signal-scaling factor

$\hat{\sigma}_o = \sqrt{S \left[1 + \frac{1}{K} \hat{\mathbf{n}}^T \hat{\mathbf{M}}^{-1} \hat{\mathbf{n}} \right]^{1/2}}$ = the adaptive filter variance normalization

These quantities are analogous to the TSCORE outputs F , S and $\hat{\sigma}$ defined for the baseline spatial filter. They are calculated for every pixel in a prefiltered clutter scene by running the adaptive spatial filter on that scene.

The smallest positive target amplitude A satisfying (33) is given by

$$A_{\min} = \frac{\left[t_o^2 \hat{\sigma}_o^2 - F^2 t_o^2 / K \right]^{1/2} - F}{S \left[1 - t_o^2 / K \right]^{1/2}} \tag{34}$$

As $K \rightarrow \infty$ the adaptive filter reduces to a precomputed matched filter, and the above solution for A_{\min} has the same form as the fixed filter TSCORE given in equation (28).

The value of A_{\min} defined in (34) is the additive peak target amplitude that is needed in the *prefiltered* scene to declare a detection at a given pixel. If the prefilter has a significant effect on the target amplitude, then this amplitude must be scaled up by an appropriate factor to get the amplitude A_{in} that would have been required in the *original* input scene. This is straightforward for a linear mean-removal filter since the peak amplitude attenuation is a constant factor for all pixels. For a nonlinear median-removal prefilter, the amplitude-dependent target attenuation for a particular scene can be estimated from curves similar to Figure 21. In this approach, the peak amplitudes calculated with the TSCORE method are mapped to input amplitudes using a least-squares polynomial relationship for A_{in} as a function of A_{out} . This relationship is found for each scene using the simulation approach described in Section 2.6.2.

To determine the threshold t_o for the adaptive filter we set $A=0$ in the detection inequality (33):

$$\frac{|F|}{\hat{\sigma}_o} \geq t_o \tag{35}$$

The TSCORE threshold is then calculated as before by computing the normalized output frame $|F|/\hat{\sigma}_o$, examining its histogram, and selecting t_o to pass no more than a specified number of false detections.

2.8.2 TSCORE Results

To compare the performance of the three filters in a consistent manner, we used the seven IRST Reference Scenes provided by NADC to measure TSCORE amplitude frames and T-90 values for an average false alarm probability of about 3×10^{-5} . These frames all have dimensions of 475x118 pixels. Excluding pixels subject to edge effects with the largest adaptive filter configuration reduces the common frame size for TSCORE computations to 440x103, for a total of 45,320 pixels. The measured TSCORE threshold settings were based on one allowed false alarm per frame, which corresponds to an average P_{fa} slightly less than the desired value. For every Reference Scene, a separate TSCORE evaluation was carried out for each of the following three filters:

- a) *GE Baseline Filter Bank*: The baseline IRST spatial filter bank described in Section 2.4.1.
- b) *Adaptive Filter 1*: A fully-adaptive filter implemented as follows:
 - 7x5/3x3 local mean removal prefilter
 - 5x3 adaptive filter window
 - 7x5 configuration of 5x3 secondary windows (34 windows)
- c) *Adaptive Filter 2*: A simplified adaptive filter based on a diagonal covariance approximation, implemented as follows:
 - 3x3/1x1 local median removal prefilter
 - 3x3 adaptive filter window
 - 7x3 configuration of 3x3 secondary windows (20 windows)

TSCORE curves of average probability of detection vs. peak input target amplitude for the three filters are shown in Figures 23, 24 and 25, respectively. The measured T-90 values are summarized in Table 5. Adaptive Filter 1 outperformed the GE baseline filter bank in 5 out of the 7 scenes, performed roughly the same in one scene, and was somewhat worse in scene 6. The average gain in target detectability for all 7 reference scenes was 2.58 dB with respect to the baseline.

The comparative performance of the simplified Adaptive Filter 2 was decidedly mixed, with improved performance obtained in 4 of the 7 scenes and degraded performance in the other 3. On an average basis, however, this filter yielded an overall gain in detection sensitivity of 1.95 dB, primarily due to the large gain obtained for scene 7.

The contrasting results obtained for Reference Scenes 6 and 7 demonstrate the relative advantages and disadvantages of a fully-adaptive detection approach. Grey-scale plots for these two scenes are shown in Figure 26.

Scene 6 contains a relatively benign sky clutter background (except for the prominent horizontal edge near the bottom of the frame). Figure 27 plots the Scene 6 TSCORE peak amplitudes for both the filter bank and Adaptive Filter 1. In these plots, "black" and "white" correspond to minimum detectable target amplitudes of 100 and zero, respectively. The darker appearance of Figure 27(b) relative to Figure 27(a) therefore indicates that the adaptive filter suffers a performance loss relative to the fixed filter bank. The precomputed "white-noise" filter in the fixed bank turns out to be closely matched to the benign clutter in most of the frame, while the adaptive filter incurs a substantial CFAR loss to match itself to the same clutter. The experiment simply confirms that a properly-matched precomputed filter will always outperform an adaptive filter, which utilizes imperfect estimates of the background statistics.

Scene 7, on the other hand, is a cloud scene containing a highly variable clutter background. Scene 7 TSCORE amplitudes for the two filtering schemes are plotted in Figure 28 on the same 100-to-zero grey-scale used for Figure 27. Here the adaptive filter significantly outperforms the filter bank, because it is able to adapt its weights more effectively to the variable clutter statistics encountered in this scene. This improved robustness is the real benefit of the fully-adaptive detection approach.

Table 6 lists the measured threshold constants used for the TSCORE computations in each reference scene, which are set to provide exactly one false alarm per frame ($P_{fa} \approx 3 \times 10^{-5}$). Pre-calculated thresholds for the adaptive filters at $P_{fa} = 3 \times 10^{-5}$ are also shown for comparison (these are read from the P_{fa} curves in Figure 6). Note that the percentage variation in the measured thresholds over the 7 scenes is




Figure 23. Cumulative Distribution of TSCORE Amplitude for Baseline Filter Bank.
 (Numbers Correspond to Reference Scenes)




Figure 24. Cumulative Distribution of TSCORE Amplitude for Adaptive Filter 1.
(Numbers Correspond to Reference Scenes)




Figure 25. Cumulative Distribution of TSCORE Amplitude for Adaptive Filter 2.
(Numbers Correspond to Reference Scenes)

Table 5. TSCORE Analysis Summary

Scene No.	Baseline Filter Bank		Adaptive Filter 1		Adaptive Filter 2	
	T90		T90	Gain	T90	Gain
1	42.5		25.0	+4.61 dB	46.0	-0.69 dB
2	30.0		25.0	+1.58	38.0	-2.05
3	45.0		36.0	+1.94	32.0	+2.96
4	76.0		53.0	+3.13	64.0	+1.49
5	33.5		36.0	-0.62	27.0	+1.87
6	21.0		34.0	-4.18	24.0	-1.16
7	78.0		42.0	+5.38	40.0	+5.80
Average Gain				+2.58 dB	+1.95 dB	





Figure 26(a) Reference Scene 6.
Reference Scene 7.





Figure 27(a) Scene 6 TSCORE Amplitude Plot for the Baseline Filter Bank.
 (b) Scene 6 TSCORE Amplitude Plot for Adaptive Filter 1.
 (Linear grey-scale with Black= +100; White=0)


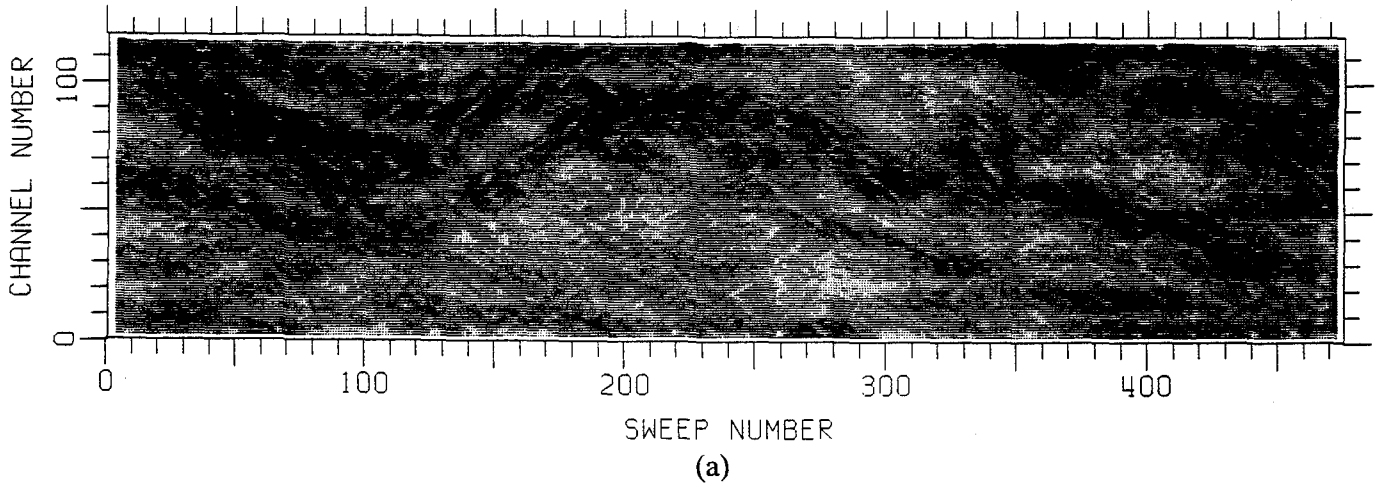


Figure 28(a) Scene 7 TSCORE Amplitude Plot for the Baseline Filter Bank.
 (b) Scene 7 TSCORE Amplitude Plot for Adaptive Filter 1.
 (Linear grey-scale with Black=+100; White=0)

Table 6. Normalized Thresholds from TSCORE Evaluation
 (Nominal $P_{fa} = 3 \times 10^{-5}$)

Scene Number	GE Baseline k	Adaptive Filter 1 t_o	Adaptive Filter 2 t_o
1	7.6	4.5	4.0
2	7.1	4.5	3.8
3	9.5	4.9	3.9
4	9.2	4.9	3.9
5	7.4	4.8	3.8
6	6.4	5.0	3.9
7	10.7	4.7	3.7
Predicted t_o ($P_{fa} = 3 \times 10^{-5}$)	—	4.5	3.9

much lower for the adaptive filters than for the baseline. Also, the measured t_0 values from the 7-scene TSCORE evaluation are remarkably close to the theoretical values of t_0 that can be predicted off-line (especially for Adaptive Filter 2). This is attributed to the CFAR property of the adaptive filter output. These results suggest that the adaptive filters would provide a more nearly constant (and more predictable) false alarm rate than the baseline processor in highly variable backgrounds.

2.9 Computation Load Comparison

An estimate of the elementary operation count per pixel was also obtained for each of the three filters evaluated above. No attempt was made to account for the relative complexity of different arithmetic operations (such as additions and multiplies), since this is a highly processor dependent factor. Our objective was to obtain a simple measure of the relative computational complexity of the adaptive filters with respect to the existing baseline.

Total computation load estimates for the three filters defined above (including prefiltering) are summarized in Table 7. As expected, Adaptive Filter 1 is significantly more complex than the baseline filter bank. This is almost entirely due to the computation associated with forming and inverting a 15x15 spatial covariance matrix estimate at each pixel. However, note that the simplified adaptive filter (Adaptive Filter 2), which does not utilize the full covariance matrix and operates with a smaller 3x3 filter window, actually requires slightly less computation than the baseline filter bank.

Table 7. Computation Load Estimates

Filter Type	Estimated Operations/Pixel
Baseline	548
Adaptive Filter 1	12093
Adaptive Filter 2	475

3.0 Conclusions of Phase I Study

The main conclusions from our Phase I study can be summarized as follows:

- 1) A fully adaptive spatial filter provides a consistent gain in target detection sensitivity but is well over an order of magnitude more complex to implement. The average performance gain for the 7 IRST Reference Scenes provided by NADC is about 2.6 dB with respect to the existing baseline filter bank. A lower T90 value (i.e., better detection performance) was obtained in 5 out of the 7 scenes.
- 2) The use of a simple nonlinear prefilter, which approximately diagonalizes the background clutter covariance matrix, allows a simplified form of the adaptive filter to be employed. Although this simplified adaptive filter is no more complex than the baseline spatial filter bank, its performance is 1.9 dB better on the average. Lower T90 values were obtained in 4 of the 7 Reference Scenes.
- 3) Compared to the baseline processor, the adaptive filter appears to provide improved false alarm regulation in unknown, variable backgrounds. This is due to the fact that the adaptive filter output is a true CFAR test statistic.
- 4) The Phase I effort was successful in demonstrating the application of adaptive spatial filtering to the IRST target detection problem, and in developing a simplified adaptive filter with acceptable detection performance. However, it is our belief that further development of adaptive algorithms which operate solely in the spatial domain cannot be justified by the relatively small performance gains obtained.

4.0 References

- [1] Hess, M.R., "Application of Clutter Scene Metrics inIRST Processor Optimization", Naval Air Development Center, Warminster PA (February 1985).
- [2] Vom Lehn, J, D.A. Langan and D.S.K. Chan, "IR Clutter Partitioning for Matched Filter Design", *SPIE Proceedings Vol. 1096: Signal and Data Processing of Small Targets* (March 1989).
- [3] Kelly, E.J., "An Adaptive Detection Algorithm", *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-22, No. 1 (March 1986).
- [4] Margalit, A., I.S. Reed and R.M. Gagliardi, "Adaptive Optical Target Detection Using Correlated Images", *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-21, No. 3 (May 1985).
- [5] Chen, J.Y. and I.S. Reed, "A Detection Algorithm for Optical Targets in Clutter", *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-23, No. 1 (January 1987).
- [6] Reed, I.S. and X. Yu, "Adaptive Multiple-Band CFAR Detection of an Optical Pattern with Unknown Spectral Distribution", to be published in *IEEE Transactions on Aerospace and Electronic Systems*.
- [7] Hunt, B.R. and T.M. Cannon, "Nonstationary Assumptions for Gaussian Models of Images", *IEEE Transactions on Systems, Man and Cybernetics*, Vol. SMC-6 (1976).
- [8] Abramowitz, M. and I.A. Stegun, *Handbook of Mathematical Functions*, Dover Publications, Inc., New York (1972).