

**Naval Undersea Warfare Center Division
Newport, Rhode Island**

**SUBSPACE DETECTION IN SUBSPACE INTERFERENCE
FOR UNDERWATER ACOUSTICS**

Walter Greene
John Impagliazzo
NUWC Division Newport

Nigel Lee
MIT Lincoln Laboratory

16 March 1999

Approved for public release; distribution is unlimited

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 16 MAR 1999		2. REPORT TYPE Technical Memo		3. DATES COVERED 16-03-1999 to 16-03-1999	
4. TITLE AND SUBTITLE Subspace Detection in Subspace Interference for Underwater Acoustics				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Walter Greene; John Impagliazzo; Nigel Lee				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Undersea Warfare Center Division,1176 Howell Street,Newport,RI,02841				8. PERFORMING ORGANIZATION REPORT NUMBER TM 990047	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research				10. SPONSOR/MONITOR'S ACRONYM(S) ONR	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES NUWC2015					
14. ABSTRACT A process called subspace detection in subspace interference (SDSI) was developed. The goal of this work was to produce an algorithm that could provide a fast estimate of the noise or interference subspace in beamformed data. The SDSI process involved separating a signal into two perpendicular subspaces, one containing noise or interference and one containing the signal alone. Detection would then be done in the subspace containing the signal only, improving detection and helping defeat countermeasures. To make the algorithm adaptive, the functions used to form the subspaces were taken from the time-frequency analysis waveforms used in the wavelet and local cosine transforms to produce good subspaces that would represent the noise well but not the signal that it is desired to detect. Projecting data into a space that is perpendicular to the noise subspace removes any of that noise in the data. The projection is performed on the matched filter for the desired signal to be detected. This should produce a modified matched filter that does not react to noise present in the data but still detects the signal. However, the output of the modified matched filter was generally not much better than that of the original matched filter. Noise couldn't be significantly separated without degrading the signal. Future work should focus on finding a better choice of bases that separate the noise and the signal better than the local cosine.					
15. SUBJECT TERMS Subspace detection in subspace interference; SDSI					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 21	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

ABSTRACT

A process called subspace detection in subspace interference (SDSI) was developed. The goal of this work was to produce an algorithm that could provide a fast estimate of the noise or interference subspace in beamformed data. The SDSI process involved separating a signal into two perpendicular subspaces, one containing noise or interference and one containing the signal alone. Detection would then be done in the subspace containing the signal only, improving detection and helping defeat countermeasures. To make the algorithm adaptive, the functions used to form the subspaces were taken from the time-frequency analysis waveforms used in the wavelet and local cosine transforms to produce good subspaces that would represent the noise well but not the signal that it is desired to detect. Projecting data into a space that is perpendicular to the noise subspace removes any of that noise in the data. The projection is performed on the matched filter for the desired signal to be detected. This should produce a modified matched filter that does not react to noise present in the data but still detects the signal. However, the output of the modified matched filter was generally not much better than that of the original matched filter. Noise couldn't be significantly separated without degrading the signal. Future work should focus on finding a better choice of bases that separate the noise and the signal better than the local cosine.

ADMINISTRATIVE INFORMATION

This memorandum was prepared under the Guidance and Control Project, principal investigator Daniel O'Neill (Code 8212). The sponsoring activity is the Office of Naval Research (ONR 333, Khine Latt).

TABLE OF CONTENTS

Section	Page
1 INTRODUCTION	1
2 BACKGROUND	1
3 TEST SETUP	4
4 RESULTS	5
5 CONCLUSIONS	12
6 REFERENCES	13

LIST OF ILLUSTRATIONS

Figure	Page
1 Detection Procedure.....	3
2 Gould Island Test Setup.....	4
3 Local Cosine Time-Frequency Display of the Coh-2 Waveform.....	5
4 Local Cosine Time-Frequency Display of the Basebanded Coh-2 Waveform.....	5
5 TF Display of Beam 3 Data Containing Structured Interference	7
6 TF Display of Beam 3 Data Containing Broadband Noise	7
7 Waveforms in the Projection Process	8
8 Results of the Technique on Data Containing Broadband Noise	8
9 Comparison of Filter Outputs on Data with Broadband Noise	9
10 Matched Filter Outputs for Various Subspace Sizes	10
11 Results of the Technique Using Second-Order Statistics	12

LIST OF TABLES

Table	Page
1. Parameters in Creating the Modified Matched Filter Fields	5

1. INTRODUCTION

The goal of this work was to produce an algorithm that could provide a fast estimate of the noise or interference subspace in beamformed data. This knowledge could then be used to remove the noise or interference, thus improving detection and helping defeat countermeasures. To make the algorithm adaptive, the functions used to form the subspaces were taken from the time-frequency (TF) analysis waveforms used in the wavelet and local cosine transforms. These waveforms can adapt to the characteristics of the data being analyzed. A good subspace would represent the noise well but not the signal that it is desired to detect. Projecting data into a space that is perpendicular to the noise subspace removes any of that noise in the data. The projection was performed on the matched filter for the desired signal to be detected. This should have produced a modified matched filter that does not react to noise present in the data but still detects the signal. Although the results were not very encouraging, there is much more work yet to be done.

2. BACKGROUND

The process developed here is called subspace detection in subspace interference (SDSI). It involves separating a signal into two perpendicular subspaces, one containing noise or interference and one containing the signal alone. Detection is then done in the subspace containing the signal only. We start with a subspace model that has the following two hypotheses:

$$\begin{aligned} H_0 : \bar{y} &= \mathbf{Q}\bar{z} + \bar{n}, \\ H_1 : \bar{y} &= \mathbf{F}\bar{c} + \mathbf{Q}\bar{z} + \bar{n}. \end{aligned} \quad (1)$$

H_0 is the hypothesis that there is no signal in subspace $\mathbf{F}$ and H_1 is the hypothesis that there is a signal present in the subspace defined by the vector $\bar{c}$. The interference subspace $\mathbf{Q}$ attempts to model the reverberation and other sources of interference. The vector $\bar{n}$ represents Gaussian-distributed random noise and is present in both cases. From the Generalized Likelihood Ratio Test (GLRT) [1], one commonly used detection statistic for this case is given by

$$t(\bar{y}) = \bar{y}^H \mathbf{P}_Q^\perp \mathbf{P}_G \mathbf{P}_Q^\perp \bar{y} = \bar{y}^H \mathbf{P}_G \bar{y}, \quad (2)$$

where

$$\mathbf{P}_Q^\perp = \mathbf{I} - \mathbf{Q}(\mathbf{Q}^H \mathbf{Q})^{-1} \mathbf{Q}^H, \quad (3)$$

and

$$\mathbf{G} = \mathbf{P}_Q^\perp \mathbf{F}, \quad (4)$$

$$\mathbf{P}_G = \mathbf{G}(\mathbf{G}^H \mathbf{G})^{-1} \mathbf{G}^H. \quad (5)$$

In these equations, $\mathbf{Q}$ is the interference subspace matrix. Each column is a basis function for the interference subspace. The goal is to find the best set of bases for the columns of $\mathbf{Q}$. The matrix $\mathbf{P}_Q^\perp$ is the projection matrix onto the space that is perpendicular to the subspace $\mathbf{Q}$. The matrix $\mathbf{G}$ is the subspace obtained from the projection of the signal subspace $\mathbf{F}$ onto the space perpendicular to the interference subspace. In this case, $\mathbf{F}$ consists of the same bases vectors used for the interference subspace $\mathbf{Q}$, so the projection is equivalent to a simple removal of bases. The matrix $\mathbf{P}_G$, defined in equation (5), is the projection matrix onto this space. If the signals and bases functions are real, the Hermitian transpose $\mathbf{H}$ can be replaced with the real transpose $\mathbf{T}$. The detection statistic $t(\bar{\mathbf{y}})$ without the projection matrix $\mathbf{P}_G$ would be a simple energy detector $\bar{\mathbf{y}}^H \bar{\mathbf{y}}$. With the projection matrix, the statistic becomes a modified energy detector where the energy is outside of the interference subspace.

By using the projection matrix on the original matched filter $\bar{\mathbf{s}}$, which is the same as the transmitted waveform, a new modified matched filter $\bar{\mathbf{s}}_m$ is created that attempts to avoid the effects of noise or interference in the received data. If $\bar{\mathbf{s}}$ is the original matched filter, detection is normally accomplished by convolving the received signal with the matched filter, $\bar{\mathbf{s}} \otimes \bar{\mathbf{y}}$, and looking for peaks in the result. It was hoped that, by using the modified matched filter $\bar{\mathbf{s}}_m \otimes \bar{\mathbf{y}}$, there would be fewer false detections due to noise or interference.

The entire procedure for detecting the target is summarized in figure 1. Starting at the box at the top left labeled “Beamformed Data,” 16,384 samples of the real beamformed data are shown. A set of data was needed that was the same length as the matched filter, so an averaging technique was performed. Averaging was accomplished by taking a windowed portion of the raw data that was 2048 points long, the same length as the matched filter. Starting at the beginning, the window was stepped through the raw data, in steps of 2048 points so that the windows didn’t overlap, and the windowed sets of data were summed together. After this was done, the results were divided by the number of windows N_w to produce the averaged data shown in the block. This procedure is given by equation (6):

$$\bar{\mathbf{y}} = \frac{1}{N_w} \sum_{i=1}^{N_w} \bar{\mathbf{y}}_i \quad (6)$$

These averaged data are the $\bar{\mathbf{y}}$ vector shown in the equations. To find the projection matrix $\mathbf{P}_G$, we need the interference matrix $\mathbf{Q}$. To find this matrix, we first selected a set of bases from which to choose the columns of $\mathbf{Q}$. In this study, the time-frequency (TF) transform used was the local cosine transform [2][3]. The columns of $\mathbf{Q}$ consist of a collection of local cosine basis functions. The goal was to select functions that represent the interference or noise well, but do not represent the desired signal well. To do this, both the averaged “noisy” data and the matched filter, which is the desired signal, were transformed using the complete local cosine transform. This produced a full set of coefficients in the local cosine domain for both signal and noise. Each coefficient

represents the component of the noise or signal that can be represented by its corresponding basis function. It was assumed that the averaged data consisted mostly of noise and interference. The signal would be present in the averaged data, but hopefully only to a small extent. The averaged data were considered to be only noise and interference. Where the coefficients for the noise are large, the corresponding bases represent the noise well. Where the signal coefficients are large, those bases represent the signal well. To determine which bases to use, the following difference was formed:

$$\bar{c} = |\bar{c}^n| - |\bar{c}^s|, \quad (7)$$

where $\bar{c}^n$ are the noise coefficients and $\bar{c}^s$ are the signal coefficients. Where the $\bar{c}$ are largest, the corresponding bases are selected to form the columns of $\mathbf{Q}$. The subspace dimension is determined by the number of bases since the local cosine bases are orthogonal to each other.

Once $\mathbf{Q}$ has been found, the projection matrix $\mathbf{P}_G$ can be determined. The matched filter was operated on by the projection matrix to produce a modified matched filter. This new SDSI-modified matched filter was similar to the original matched filter but was perpendicular to the noise. The detection process consisted of convolving the matched filter with the beamformed data and thresholding the output. Large values in the output of the matched filter are called detections and should, but don't always, indicate the presence of the signal. To test the ability of the modified matched filter to remove noise and interference, it was convolved with the original beamformed data to give the SDSI output. This output was compared with the convolution of the original matched filter with the same data. The hope was that by making the matched filter less sensitive to the noise in the beamformed data, there would be fewer false detections.




Figure 1. Detection Procedure

3. TEST SETUP

To test the capabilities of the modified matched filter, real data collected from the Gould Island experiments run on 19 - 28 May 1997 were used. In these experiments, a Torpedo Mk 48 ADCAP nose section was mounted on the East Elevator on Gould Island and lowered to a depth of 30 feet. This test was a one-way transmission where the target returns were generated via a synthesized waveform generator and transmitted from a BQR-7 transducer located in 60 feet of water and down range approximately 100 yards. In this way, the ADCAP was a receiver only and the signals transmitted were contaminated by noise in only one direction. A countermeasure device was suspended from a range craft and positioned approximately 180 yards down range. This device provided two different types of interference, structured interference and broadband noise. The data received at the ADCAP were beamformed to form nine different beams at various angles. Figure 2 shows the test setup with the angles of the target generator, the interference, and the beams. As figure 2 shows, the strength of the interference in the beamformed data depended on the beam number. Beams 2 and 3 showed the strongest interference and beams 5 and 6 had the strongest target.




Figure 2. Gould Island Test Setup

The waveform transmitted was a linear stepped FM called Coh-2. It was sampled to give the 2048 data points, along with its time-frequency image, shown in figure 3. Initially, because of the need to form matrices such as $\mathbf{P}_G$, which would be a 2048 by 2048 matrix, the data were basebanded and subsampled to produce a 256-point complex waveform. The real part of the basebanded waveform is shown in figure 4. In later tests, the transform was done directly, allowing the original waveforms to be used. The results given in this memorandum are all from the latter results using the full original 2048-point waveform.




Figure 3. Local Cosine Time-Frequency Display of the Coh-2 Waveform




Figure 4. Local Cosine Time-Frequency Display of the Basebanded Coh-2 Waveform

4. RESULTS

The results when applying the techniques described in section 2 on the data were not as good as desired. The output of the modified matched filter was, in general, not much better than the output of the original matched filter. This section presents some of the results using various combinations of data and parameters.

There were several parameters (shown in table 1) that could be adjusted to affect the form of the modified matched filter.

Table 1. Parameters in Creating the Modified Matched Filter Fields

Parameter	Description
Subspace Size	The number of columns in the interference matrix $\mathbf{Q}$.
Type of Bases	For this study either the local cosine or the local sine basis functions were used.
Number of Windows	This is the number of time intervals into which the signal being transformed is split. It determines the resolution of the transform in both time and frequency.
Length of Window Overlap	In the local cosine and local sine transforms, adjacent windows overlap, yet the bases remain orthogonal. This determines the length of the overlap. It was not adjusted in this study and was kept at one-half of the window length.

The subspace size is probably the most critical in this study since the types of bases used were limited to only local cosine and local sine bases. Using the sampled data, which have a length of 2048 points, there are a total of 2048 functions in the complete local cosine basis for the space. It was desired to find a small subset of these basis functions to represent the noise/interference. Since the matched filter is projected onto the space perpendicular to the subspace represented by the selected bases, the smaller the number of bases used, the less the signal is affected. The number of windows used should, in general, be matched to the characteristics of the waveform being transformed. As it turned out, for the waveform with 2048 points, 64 windows were a good match, although this number should be further explored. Since the length of the window overlap is not a critical parameter, a value of $\frac{1}{2}$ was always used.

Figure 5 shows the time-frequency display of one set of beamformed data. This is beam 3 data with structured interference present. The size of the windows shown (the horizontal length of the boxes) gives an idea of the resolution in time and frequency used. Since the matched filter is 2048 points long and is divided into 64 windows of 32 points each, the entire beamformed data, containing 16,384 points, is split into 512 windows of 32 points each. The transmitted signal is present starting at about location 2250. The beamformed data contain a lot of noise at frequencies over a large bandwidth along with several tones, shown as lines in the TF display. Figure 6 is the time-frequency display for beam 3 of the beamformed data that contain only broadband noise. The transmitted signal is present also at approximately location 2250. Figure 7 shows many of the waveforms involved in the detection algorithm. All of these were generated while processing the beamformed data containing structured interference shown in figure 5. Figure 7a is the TF pattern of the original matched filter, and figure 7b is the TF pattern of the averaged data. Each box in the TF display corresponds to one coefficient in the local cosine transform. The matched filter is an FM sweep, while the averaged data contain mostly noise, although the signal is present in it. By taking the difference between the signal shown in figure 7a and the noise shown in figure 7b, we find where the noise is large and the signal is not. Figure 7c shows the projection of the averaged data onto the space defined by the largest 1024 components of this difference. It can be seen that the locations where the signal is strong are not included in the space. The matched filter, shown in figure 7a, is then projected into the space perpendicular to that shown in figure 7c. The resulting modified matched filter is shown in figure 7d. The original matched filter and the new modified matched filter waveforms are shown in figure 7e. Finally, the outputs of the correlation of the two matched filters with the original beamformed data are shown in figure 7f. Figure 7f demonstrates that the technique does not work as well as desired. Even with half of the noise components removed from the matched filter, it still correlates with the noise in the data starting at around 12,000. Also, there was no improvement in the peak correlation with the true target.

Figure 7 involved a fairly easy example since even the original matched filter shows a clear peak at the expected location. Figure 8 demonstrates the results on a more difficult problem. In this case, the data are from beam 3 and contain broadband noise, as shown in figure 6. It should be noted that, for both the original matched filter and the

modified matched filter, there is not a significant peak near location 2250. Also, there is no improvement in reducing the other peaks that are caused by noise. Varying the beams changes the difficulty of the problem. Figure 9 shows the output of the original matched filter and the modified matched filter for all nine beams of the data with broadband noise. Only the beams pointed toward the transmitter and away from the interference source show a clear peak.




Figure 5. TF Display of Beam 3 Data Containing Structured Interference




Figure 6. TF Display of Beam 3 Data Containing Broadband Noise




Figure 7. Waveforms in the Projection Process




Figure 8. Results of the Technique on Data Containing Broadband Noise

Another variable that was adjusted was the size of the subspace used to represent the noise/interference. The matched filter was projected away from this subspace so that a smaller subspace affected the matched filter less. In this case, a total of 2048 bases could be used to represent the noise. Using no bases left the matched filter unchanged, while using all 2048 bases zeroed out the matched filter. Figure 10 shows the results for various sizes of subspace for the problem with broadband noise in beam 3. The first row is the original matched filter and its output. The other rows represent various size subspaces and should be compared with the first row. It seems that the size of the subspace does not make much difference in increasing the size of the peak at 2250 relative to the other false peaks.




Figure 9. Comparison of Filter Outputs on Data with Broadband Noise




Figure 10. Matched Filter Outputs for Various Subspace Sizes

Probably the most likely area for error in the technique described so far is in determining the averaged data. Averaging the data is inherently error-prone. There is no reason to believe that the statistics of the resulting averaged data represent the actual statistics of the beamformed data in general. While this method looked at first-order statistics of the noise, i.e., averages, the following equations determine a second-order method of finding the best basis functions to represent the noise. If $\bar{y}_i$ is the i^{th} windowed subset of the input beamformed data, we can write the problem as the following minimization problem:

$$\min_{\{\bar{c}_i\}_{i=1}^{N_w} \& \mathbf{Q}} \frac{1}{N_w} \sum_{i=1}^{N_w} \|\bar{y}_i - \mathbf{Q}\bar{c}_i\|^2. \quad (8)$$

We want the best set of basis functions from the set $\mathbf{Q}$, and the best coefficients $\bar{c}_i$ for each of the N_w windows to represent the data. First minimize over $\bar{c}_i$ for a given $\mathbf{Q}$ and i to get the optimum $\hat{c}_i$:

$$\hat{c}_i = (\mathbf{Q}^T \mathbf{Q})^{-1} \mathbf{Q}^T \bar{y}_i. \quad (9)$$

Substituting this back into equation (8), we get

$$\frac{1}{N_w} \sum_{i=1}^{N_w} \|\bar{y}_i - \mathbf{P}_Q \bar{y}_i\|^2. \quad (10)$$

where

$$\mathbf{P}_Q = \mathbf{Q}(\mathbf{Q}^T \mathbf{Q})^{-1} \mathbf{Q}^T \quad (11)$$

is the projection operator. Taking the square, we get

$$\begin{aligned} & \frac{1}{N_w} \sum_{i=1}^{N_w} \bar{y}_i^T \bar{y}_i - \frac{2}{N_w} \sum_{i=1}^{N_w} \bar{y}_i^T \mathbf{P}_Q \bar{y}_i + \frac{1}{N_w} \sum_{i=1}^{N_w} \bar{y}_i^T \mathbf{P}_Q^T \mathbf{P}_Q \bar{y}_i \\ &= \frac{1}{N_w} \sum_{i=1}^{N_w} \bar{y}_i^T \bar{y}_i - \frac{1}{N_w} \sum_{i=1}^{N_w} \bar{y}_i^T \mathbf{P}_Q \bar{y}_i. \end{aligned} \quad (12)$$

Since we are minimizing over $\mathbf{Q}$, there is only one term that includes $\mathbf{Q}$ and it has a negative sign in front, so minimizing this sum is equivalent to maximizing

$$\frac{1}{N_w} \sum_{i=1}^{N_w} \bar{y}_i^T \mathbf{P}_Q \bar{y}_i, \quad (13)$$

which is the same as

$$\text{trace} \left\{ \mathbf{P}_Q \frac{1}{N_w} \sum_{i=1}^{N_w} \bar{y}_i \bar{y}_i^T \right\}. \quad (14)$$

If $\hat{\mathbf{R}}_y$ is the correlation matrix given by

$$\hat{\mathbf{R}}_y = \frac{1}{N_w} \sum_{i=1}^{N_w} \bar{y}_i \bar{y}_i^T, \quad (15)$$

equation (14) becomes

$$\text{trace}\{\mathbf{P}_Q \hat{\mathbf{R}}_y\} = \text{trace}\left\{\sum_{i=1}^{N_w} \bar{q}_i \bar{q}_i^T \hat{\mathbf{R}}_y\right\}. \quad (16)$$

Here $\bar{q}_i$ are the columns of the projection matrix $\mathbf{P}_Q$. Because the basis functions are orthogonal, the $\bar{q}_i$ are also the columns of $\mathbf{Q}$. The final result is that we want to maximize

$$\max_{\bar{q}_i} \sum_{i=1}^N \bar{q}_i \hat{\mathbf{R}}_y \bar{q}_i^T. \quad (17)$$

Here, N is the desired dimension of the subspace. In the algorithm, the matrix-vector products shown in equation (17) are calculated for all the basis vectors and the top N are selected.

The results using this new selection technique are shown in figure 11. There was not much difference between these results and the ones shown in figure 8. The same data were used in each case, and the only difference was the method of selecting the 1024 basis functions that make up the columns of $\mathbf{Q}$.




Figure 11. Results of the Technique Using Second-Order Statistics

5. CONCLUSIONS

The theory behind the work presented here is good, but the structure of the problem may be what leads to a non-optimal solution. The noise and interference present in this problem were fairly wideband and exist throughout the time of interest, especially in the broadband noise case. Because of this, it takes a large number of coefficients to represent the noise. If enough coefficients are taken to represent the noise, then the signal is degraded in the process. Also, significant noise exists in the same time/frequency bins as the signal, so it cannot be removed without removing the signal. This problem might be lessened with a different choice of bases that separate the noise and signal better than the local cosine. Future work should focus on finding a better choice of bases to separate the noise and the signal.

6. REFERENCES

1. Steven M. Kay, *Fundamentals of Statistical Signal Processing, Volume II, Detection Theory*, Prentice-Hall PTR, 1998.
2. Walter Greene, *The Local Trigonometric Transforms and Their Use in Solving Partial Differential Equations*, Ph.D. Thesis, Brown University, September 1998.
3. Gil Aharoni, Amir Averbuch, Ronald Coifman, and Moshe Israeli, "Local Cosine Transform – A Method for the Reduction of the Blocking Effect in JPEG," *Journal of Mathematical Imaging and Vision*, March 1993, pages 7-38.

DISTRIBUTION LIST

External:

Office of Naval Research (K. Latt (ONR 333))
MIT Lincoln Laboratory (N. Lee)

Internal

Codes: 5441 (2)
821 (J. Kelly)
8212 (J. Fay, W. Greene, J. Impagliazzo (5), M. Medeiros, A. Mirken,
and D. O'Neill)

Total: 15

