

LIBRARY USE ONLY

5441
TM97-2044

NUWC-NPT Technical Memorandum 972044

Copy 1

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



**WAVELET-BASED DATA COMPRESSION FOR COMMUNICATION OF
SIDE SCAN SONAR IMAGES**

J. Impagliazzo and W. Greene
Weapons Technology and Tactical Vehicle Systems Department

G. Dobeck, Q. Huynh, and L. Smedley
Naval Surface Warfare Center, Dahlgren Division, Coastal System Station
Panama City, Florida

B. Jawerth
Summus Ltd. and University of South Carolina

Leonard Mygatt
Summus Ltd.

15 April 1997

UNCLASSIFIED
NAVAL UNDERSEA WARFARE CENTER
DIVISION NEWPORT
NEWPORT, RHODE ISLAND 02841-1700
RETURN TO: TECHNICAL LIBRARY

Approved for public release; distribution is unlimited.

LIBRARY USE ONLY

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 15 APR 1997	2. REPORT TYPE Technical Memo	3. DATES COVERED 15-04-1997 to 15-04-1997			
4. TITLE AND SUBTITLE Wavelet-Based Data Compression for Communication of Side Scan Sonar Images		5a. CONTRACT NUMBER			
		5b. GRANT NUMBER			
		5c. PROGRAM ELEMENT NUMBER			
6. AUTHOR(S) J. Impagliazzo; W. Greene; G. Dobeck; Q. Huynh; L. Smedley		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 972044			
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research		10. SPONSOR/MONITOR'S ACRONYM(S) ONR-321TS			
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES NUWC2015 Additional authors: B. Jawerth; Leonard Mygatt					
14. ABSTRACT Both the Remote Mine-Hunting System (RMS) and the Long-Term Mine Reconnaissance System (LMRS) will deploy remote mine countermeasure (MCM) vehicles with side scan sonar as a principal sensor. Wavelet-based data compression techniques are presented for preprocessing side scan sonar images communicated from remotely deployed MCM vehicles to the host platforms through a single communication channel. To satisfy bandwidth, signal-to-noise ratio, and real-time processing requirements, orthogonal and biorthogonal wavelet bases have been evaluated for algorithms achieving image compression ratios at 25:1, 50:1, and 100:1. These compression ratios are required while preserving the visual clues human operators use to classify mine-like objects and computational clues used by automated classification algorithms. Fleet side scan sonar images have been processed and signal degradation evaluated based on comparative performances of an automated classification algorithm. Results from testing on compressed and uncompressed images are presented for the automated detection and classification algorithm developed under the Office of Naval Research sponsorship.					
15. SUBJECT TERMS Remote Mine-Hunting System; RMS; Long-Term Mine Reconnaissance System; LMRS; side scan sonar					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 40	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

ABSTRACT

Both the Remote Mine-Hunting System (RMS) and the Long-Term Mine Reconnaissance System (LMRS) will deploy remote mine countermeasure (MCM) vehicles with side scan sonar as a principal sensor. Wavelet-based data compression techniques are presented for preprocessing side scan sonar images communicated from remotely deployed MCM vehicles to the host platforms through a single communication channel. To satisfy bandwidth, signal-to-noise ratio, and real-time processing requirements, orthogonal and biorthogonal wavelet bases have been evaluated for algorithms achieving image compression ratios at 25:1, 50:1, and 100:1. These compression ratios are required while preserving the visual clues human operators use to classify mine-like objects and computational clues used by automated classification algorithms. Fleet side scan sonar images have been processed and signal degradation evaluated based on comparative performances of an automated classification algorithm. Results from testing on compressed and uncompressed images are presented for the automated detection and classification algorithm developed under the Office of Naval Research sponsorship.

ADMINISTRATIVE INFORMATION

The sponsoring activity for the work reported in this technical memorandum was the Office of Naval Research (ONR-321TS (W. Ching and R. Jacobson)).

1. INTRODUCTION

Currently the Remote Mine-Hunting System (RMS) program is developing a remotely deployed mine-hunting system using a snorkeling unmanned underwater vehicle (UUV) that tows a side scan sonar fish. With 15.24-cm (6-in.) resolution in the direction of travel at speeds in advance of 6 knots, a port and starboard side scan sonar image with 1024-by 511- 8-bit pixels can be collected every 25 seconds. Without data compression, it could take as long as 45 minutes to transmit 25 seconds of data over a standard radio frequency (RF) or satellite link. The actual time required for lossless transmission depends on the encoding scheme allowed by the signal-to-noise ratio (SNR) of the channel.

When the sonar pings, an underwater mine or similar object moored or sitting on the ocean bottom prevents sound from a side scan sonar system from reaching the seafloor for some distance beyond the object. This occurrence produces a characteristic highlight and adjacent shadow highly localized in the side scan sonar image. An automatic target recognition algorithm, therefore, might categorize the dimensions and relative intensity of the highlight and shadow to determine if the object should be classified as mine-like. This algorithm would, however, be susceptible to false alarms if the data compression techniques employed produced artifacts in the compressed image resembling the highlight-shadow characteristic. While the false alarm rate must be minimized, a minimum image compression ratio on the order of 25:1 to 100:1 is (1) required to keep pace with the real-time data acquisition rate of a remotely deployed vehicle and (2) necessary to satisfy bandwidth and SNR limitations of acoustic, high-frequency (HF), or satellite communications links used to communicate the images to a host platform for analysis.

The Navy Imaging Database at the Naval Surface Warfare Center, Dahlgren Division (NSWCDD), Coastal System Station (CSS), was used in the research described in this technical memorandum (TM). The images in this database were produced by a fleet side scan sonar towed by a helicopter. The database consists of 60 images, 30 of which have been designated as training images, and 30 as testing images. Fifteen of the 30 training images contain one mine signature each; 16 of the 30 testing images contain one mine, and one testing image contains two mines, for a total of 33 mine signatures. The mines in this database are cylindrical bottom mines that typically have both a highlight and shadow signature. For this sonar, a typical mine signature has around 36 pixels (picture elements) in the highlight region and about 120 pixels over the shadow zone, but these amounts vary greatly. The data for each image consist of a matrix of 1024- by 511- 8-bit unsigned integers. For processing purposes, the last column is duplicated to give 512 columns.

Figure 1, image si000206 (sonar image number 206), is an example from the Navy Imaging Database. Near range is at the top of the figure, and far range is at the bottom. Cross-range (the direction of travel) is horizontal across the image. The near range appears to be smooth while the far range is rough; abnormalities appear as striations in the last 5-percent of range. The apparent smoothness of the near range is due to the higher angle of sound incidence. In reality, the roughness and tracks are distributed uniformly over the image. The dark tracks, in many cases, are caused by fishermen dragging shrimp nets. The orientation of the tracks is also evenly distributed, but more horizontal tracks show up because of their acoustic shadow. The axes shown are in meters. Resolution in the direction of travel, approximately 15.24 cm, is a function of the speed of the tow vehicle and the round trip time of the ping to the farthest range. Resolution in range, also approximately 15.24 cm, is a function of the number of beams and the maximum range. The minimum resolution in range is also determined by the size of objects that must be detected. Two mines are found at locations (65, 56) and (131, 24) and are shown in the blow ups in figure 1. The first mine, at (65, 56), has a modest horizontal highlight and a

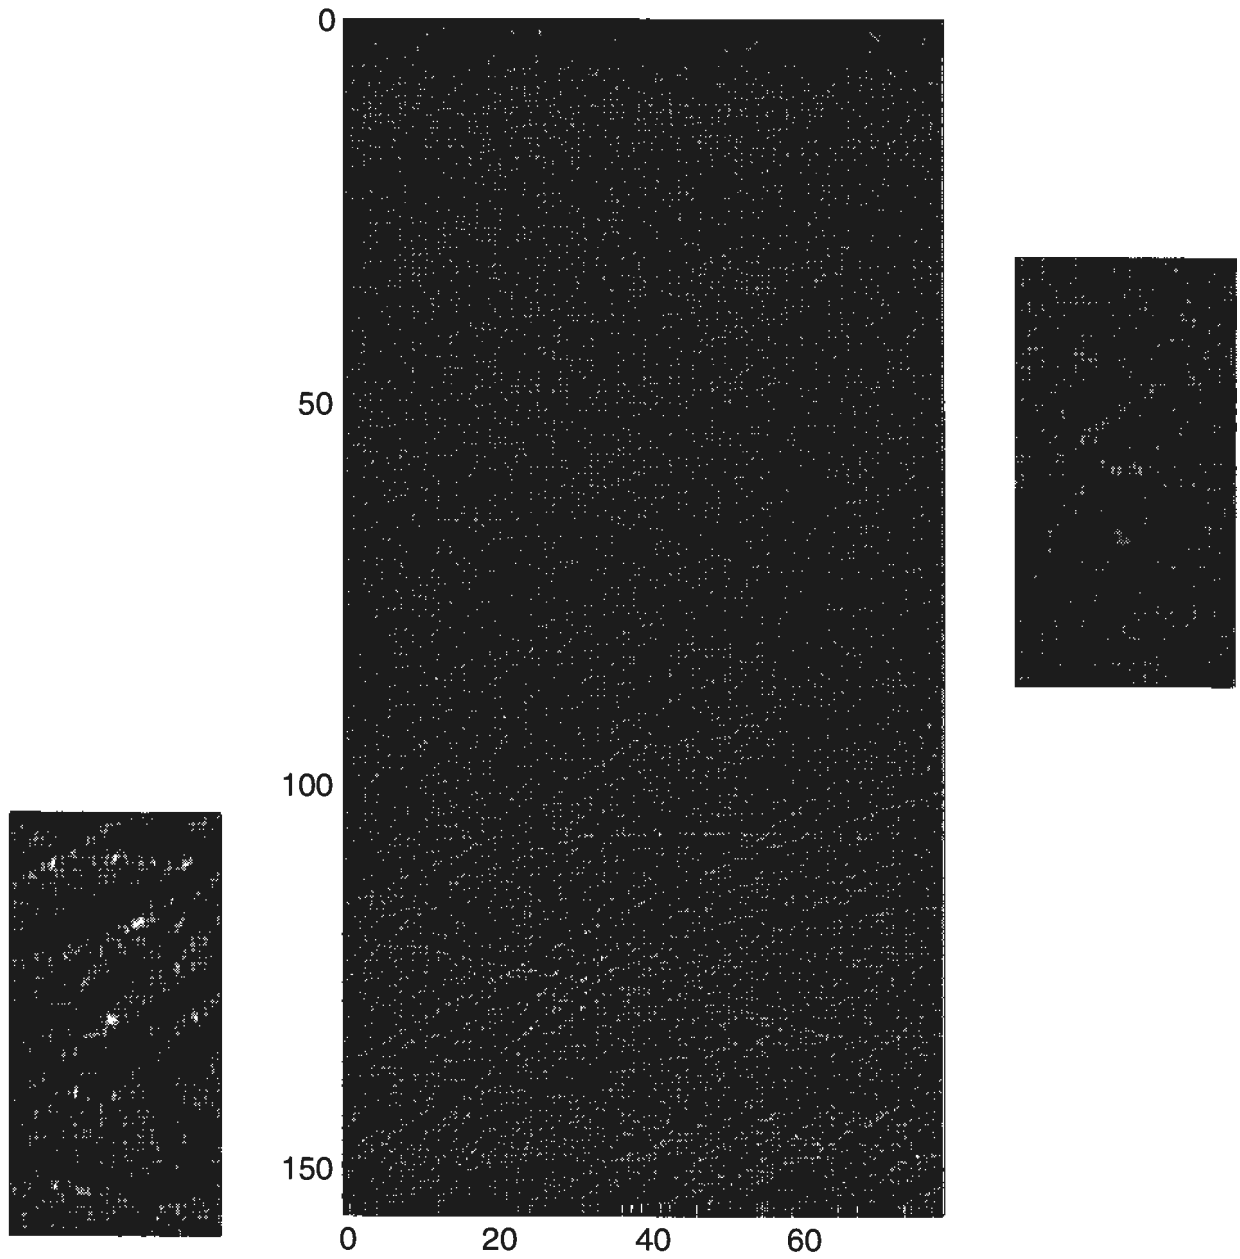


Figure 1. Sample Fleet Side San Image, File si000206, with Blow Up of Mines at (131, 24) (Left) and (65, 56) (Right)

pronounced shadow. The second mine, at (131, 24), is difficult to see; it has a small, strong highlight and a small shadow, which is somewhat disguised because it is located on the edge of a track.

2. BACKGROUND

IMAGE COMPRESSION ALGORITHM

A lossless image compression algorithm is one that allows perfect reconstruction of the original image. NUWC's transform-based image compression algorithm is a lossy algorithm—an algorithm that allows reconstruction of the image, but with some degradation in image fidelity. This particular lossy algorithm consists of three primary components: (1) transform, (2) quantizer, and (3) encoder.

A transform is an invertible process that provides optimal or near optimal identification of critical image data. A quantizer grades the value of the image information provided by the transform and selects the information most “valuable” to minimizing error measure for the desired compression level. An encoder losslessly encodes the information remaining after quantization in order to remove redundancy in the data and, as a result, decreases the number of bits required to represent the transformed and quantized image information.

The losses occur mainly in the quantizer as the result of data roundoff or loss of data not considered valuable when compressing. The inverse operations of these three component processes occur during decompression (reconstruction) of the compressed image data.

DESIGN CONSIDERATIONS

The key to a successful lossy compression algorithm is also the key to a successful analysis algorithm: *efficient identification of information meaningful to the end user*. In the case of compression, it is important to identify this information so that it is preserved through any compression process. In the case of analysis, it is important to identify this information so that only the valuable information is passed on to the following detection/classification processes. In practical applications, the selection algorithm must consider several real constraints in arriving at an “optimal” solution.

High Compression with Good Image Quality

The compression algorithm should seek to maximize the preservation of the image information identified as important by the end user, a goal often overlooked in the application of image compression algorithms. Many times a simple comparison measure, e.g., peak signal-to-noise ratio (PSNR), is not sufficient to distinguish the advantages of one approach over another in maximizing mission success. Performance comparisons between algorithms must ultimately be measured in relation to the end user success criterion.

Another important requirement is that the compression and decompression processes be designed to complement each other in their response to lost information. For example, the forward and inverse transforms are invertible pairs when there is no loss of information. How the forward and inverse pair react to the introduction of the nonlinear quantization error can have a major impact on performance. Improper selection of a transform pair may result in significant distortions in the reconstructed image as well as the inability of the user to correctly identify the desired information.

Graceful Degradation with Increasing Compression

The ability of an algorithm to provide useful information over a broad range of compression ratios provides a solid foundation for hierarchical approaches to image observation. For example, image compression can be set at a very high level until “something” of coarse interest is identified, after which compression levels can be reduced to provide finer detail information to the observer. However, if an algorithm has a catastrophic failure point, then these kinds of operations are not possible. Also, this characteristic ensures that small changes in compression level will not result in dramatic changes in image quality.

Broadly Optimal for Many Image Sources

Because the precise characteristics of an image cannot be known before compression, it is important that the algorithm be able to achieve similar performance over a broad range of image statistics. It is easy to “tune” an algorithm to a specific test data set, but what happens if the real data vary in even a marginally significant way from the test data? Unless the application has sufficiently narrowed down the uncertainty in the source data, there exists the possibility of catastrophic failure. For practical applications, an adaptive algorithm or one that is broadly optimal is required.

Error Resilient

Because the compressed data must ultimately be stored or transmitted through some channel medium, the source of channel errors should be minimized. In many cases it is not possible to minimize the error, so the data must be protected. The classic approach to encoding is to treat the channel and source encoding tasks as separate tasks. This approach assumes equal information weighting of all the data in the source, yet the information weighting of the data is most often not uniform. Consequently, severe losses in the compression efficiency result because all data must be protected to the same extent as the least important data, particularly for noisy channels. A proper analysis of the data information weighting can usually lead to efficient classifications of the data into information weight sets. An appropriate level of channel coding may be applied to each set to protect each set according to its total information content. The result is an overall improvement in compression efficiency.

Fast, Symmetric, Low Latency

Being fast, symmetric, and having low latency are three important, but often overlooked, constraints in image compression algorithm development. First, given the time, an algorithm can be designed that will produce very high compression efficiency. However, practical considerations limit the time allowed to compress (or decompress) an image because there is always a tradeoff between processing speed and compression efficiency. Second, because different applications may not provide equal processing ability to both encoder and decoder, the algorithm complexity should be approximately equal for both compression and decompression. This equal complexity permits flexible design considerations while maintaining compressed image file compatibility. Third, the algorithm should be able to process the information as soon as information is available. Long delays between data input and output can lead to buffer and memory management problems.

Progressive Transmission

The ability to provide progressive transmission adds additional flexibility, particularly in bandwidth-constrained environments. Specifically, progressive transmission allows rapid identification of critical information by transmitting “important” information first and then presenting this information in a format that indicates the benefit of transmitting finer detail

information. The basic architecture required for progressive transmission also ensures a hierarchical redundancy to the data. This redundancy may be used to efficiently match the transmitted data to the transmission channel bandwidth or to match a single transmission to several different channel bandwidths.

Built-In Features and Enhancements

Certain image enhancements can greatly improve mission success. The ability of a compression algorithm to perform these enhancements with low additional computational overhead provides a means to efficiently normalize the data in some way or to efficiently preserve certain image features through the lossy compression process. Computing the enhancements within the compression process controls the effect of compression losses on the enhanced data.

Simple Software or Hardware Design—Low Complexity

When possible, the compression algorithm should be designed to match the capabilities of the processing hardware—microprocessor or application-specific integrated circuit (ASIC). Maintaining integer operations, binary divides, etc., in the algorithm design greatly enhances the speed and reduces the error propagation of the algorithm.

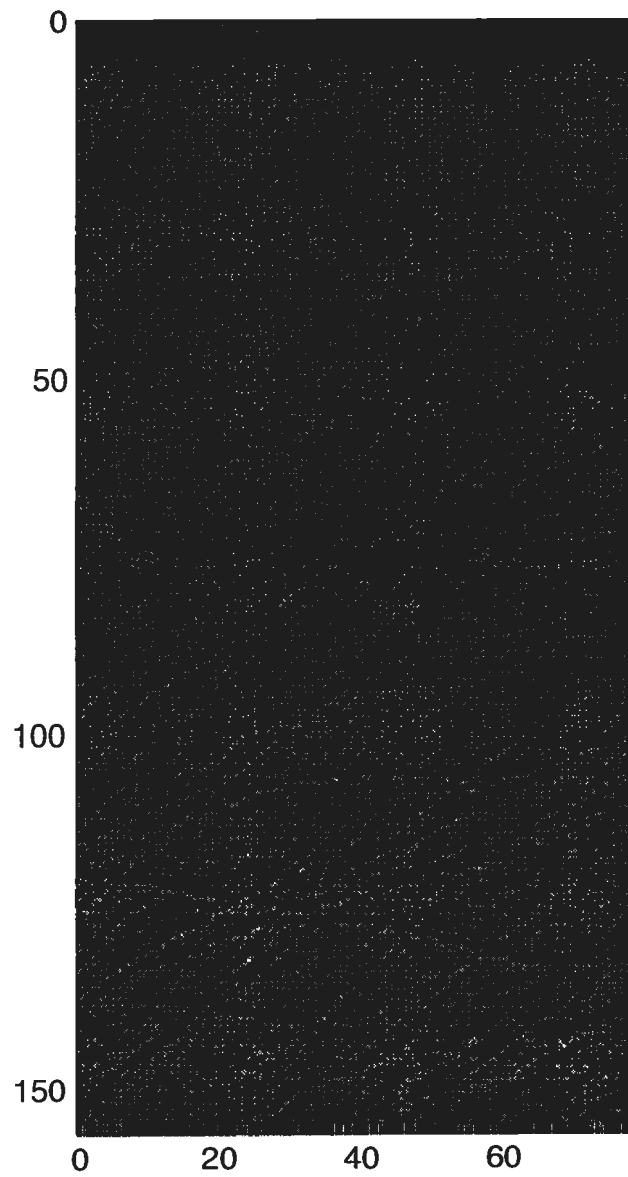
Low Power—Low Memory

In practical, embedded applications, memory and power are primary concerns. The compression algorithm should be able to easily work in limited power and memory environments without loss in performance. For example, with wavelet-based image compression, the typical approach is to first transform the entire image data and then start subsequent quantization and encoding. This approach requires a large “scratch-pad” memory to hold the transform data during calculation of the transform coefficients. For certain types of wavelet transforms, efficient computational organizations that significantly reduce the required scratch-pad memory exist.

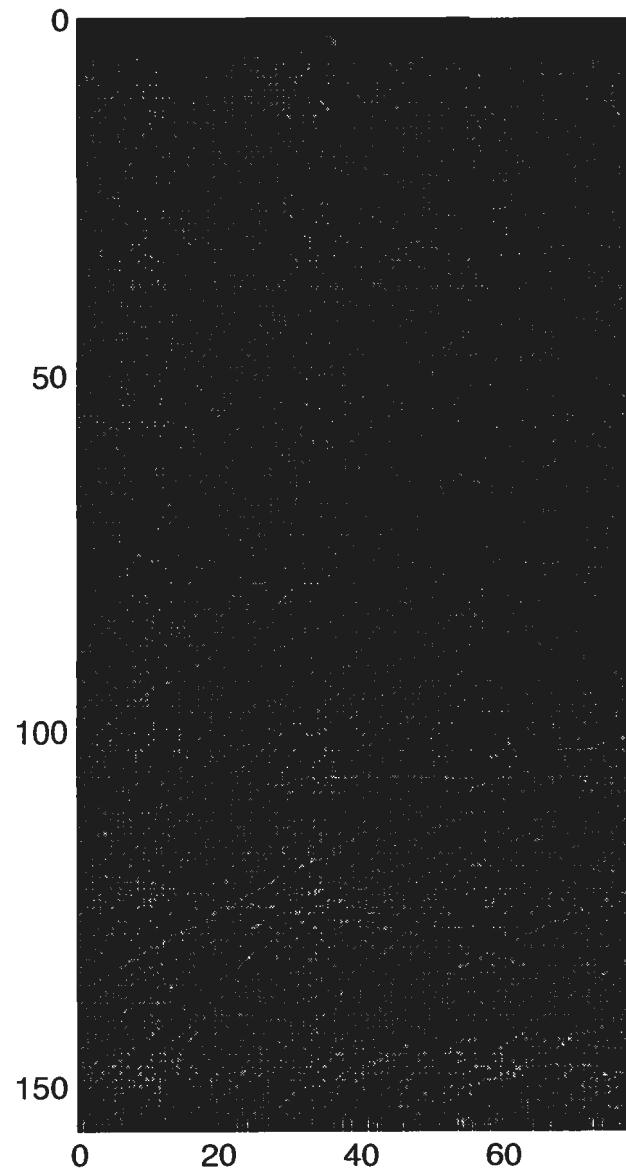
JOINT PHOTOGRAPHIC EXPERT GROUP (JPEG) TECHNIQUE

The standard technique used for still image compression is referred to as “the Joint Photographic Expert Group (JPEG) technique” (reference 1). The source image is divided into 8-by 8-pixel blocks, and the discrete cosine transform (DCT) is computed for the 64 pixels producing an average, or 0 frequency, value and 63 frequency values. Depending on the degree of compression required, these 64 values are quantized to produce a minimum set representing each block. Because at least one value is required to represent each block, the maximum compression ratio achievable is 64:1. The theoretical limit of the DCT technique, without regard to image quality, is a compression ratio of 64:1. The reconstructed image, therefore, consists of blocks assigned the average gray level of the enclosed 8 x 8 pixels. Lossless compression techniques are then applied to the composite of pixel blocks. Using the JPEG algorithm, the original and reconstructed images for compression ratios of 25:1 and 50:1 were generated for file si000206 and are shown in figure 2. The regions containing each of the two targets are magnified and shown in figure 3 and figure 4.

An image can also be decomposed with well-localized orthogonal and biorthogonal basis functions. The compression ratios and image quality possible with a wavelet-based data compression algorithm are significantly better because the wavelet basis function is orthogonal and well localized in time and scale and can, therefore, capture substantially more information per coefficient than the DCT (reference 2).



*Figure 2a. JPEG Compression of Image si000206
(Original Image)*



*Figure 2b. JPEG Compression of Image si000206
(25:1 Compression Ratio)*

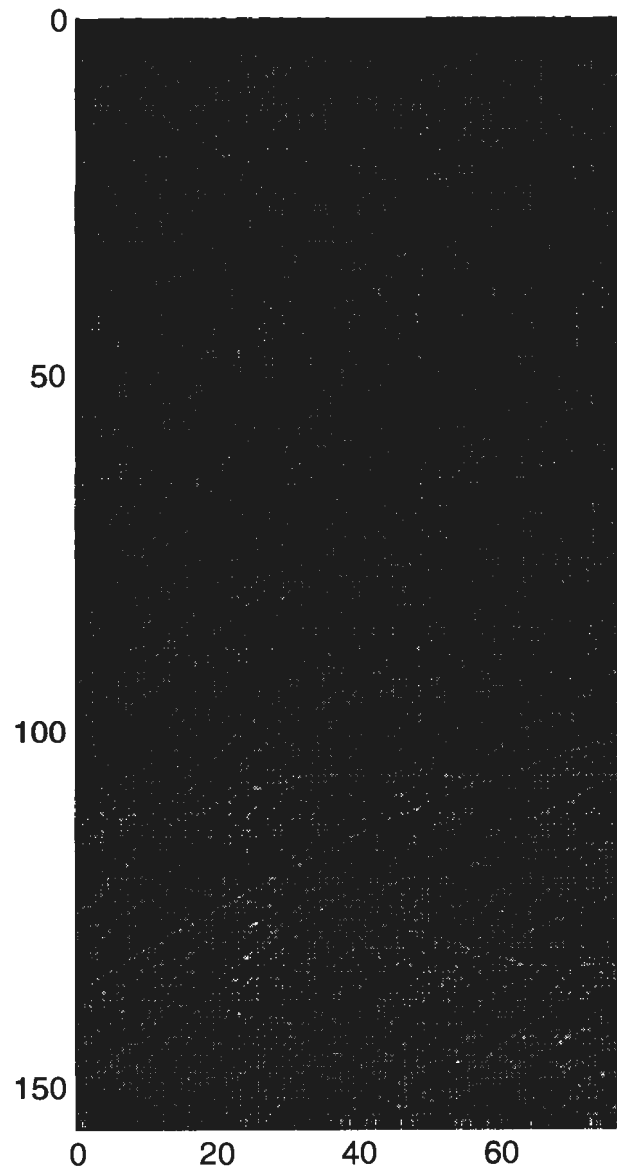


Figure 2c. JPEG Compression of Image si000206 (50:1 Compression Ratio)

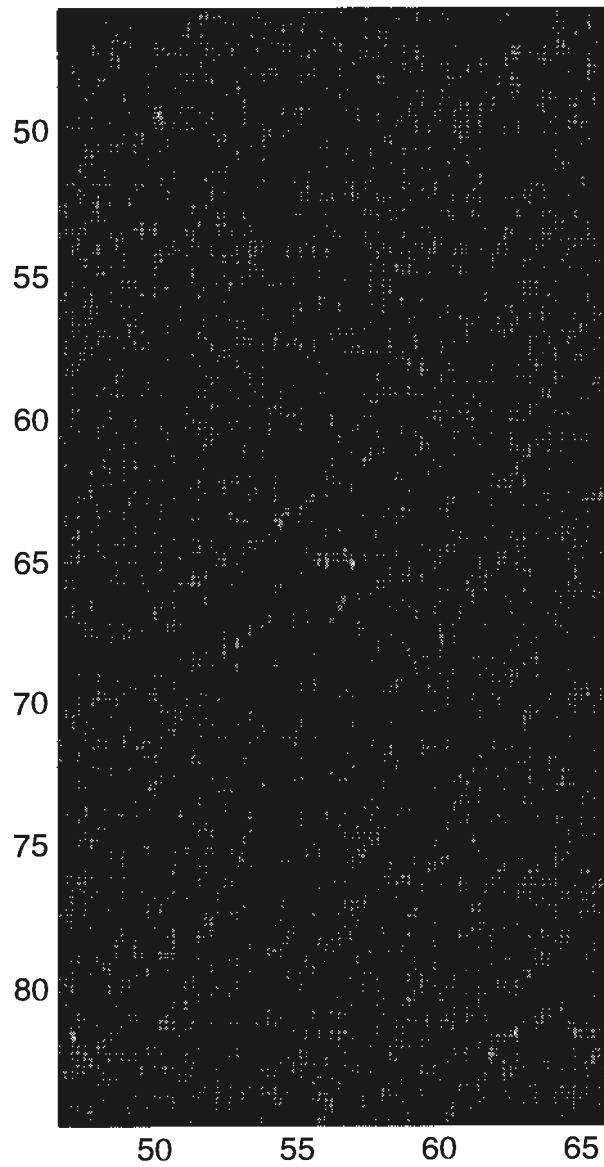


Figure 3a. JPEG Compression of Image si000206 Showing Target Location (65, 56) Magnified (Original Image)

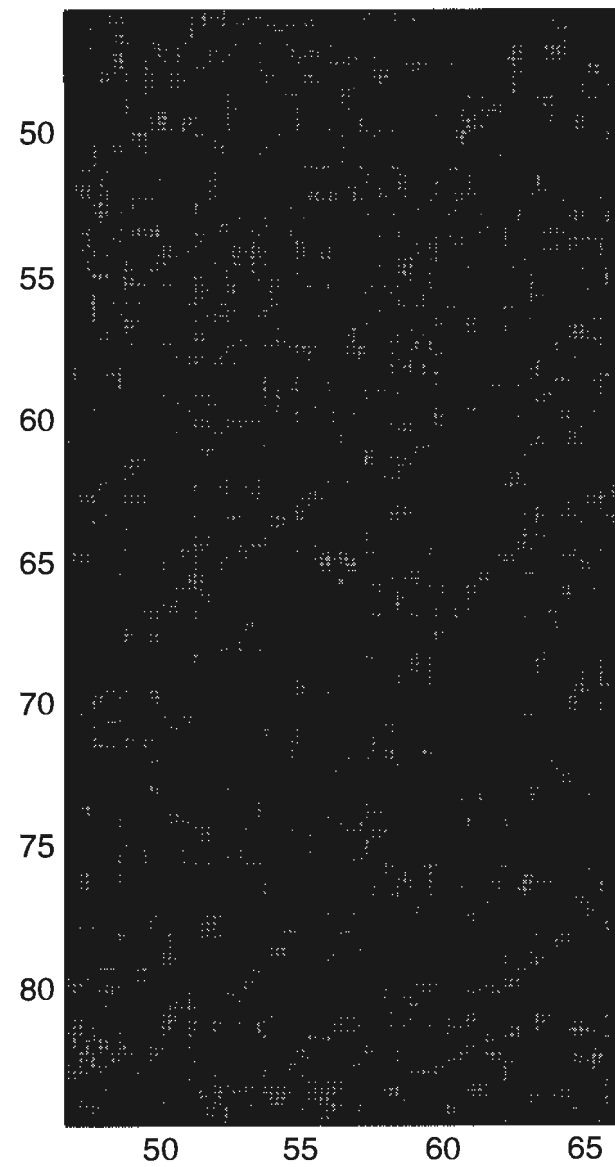


Figure 3b. JPEG Compression of Image si000206 Showing Target Location (65, 56) Magnified (25:1 Compression Ratio)

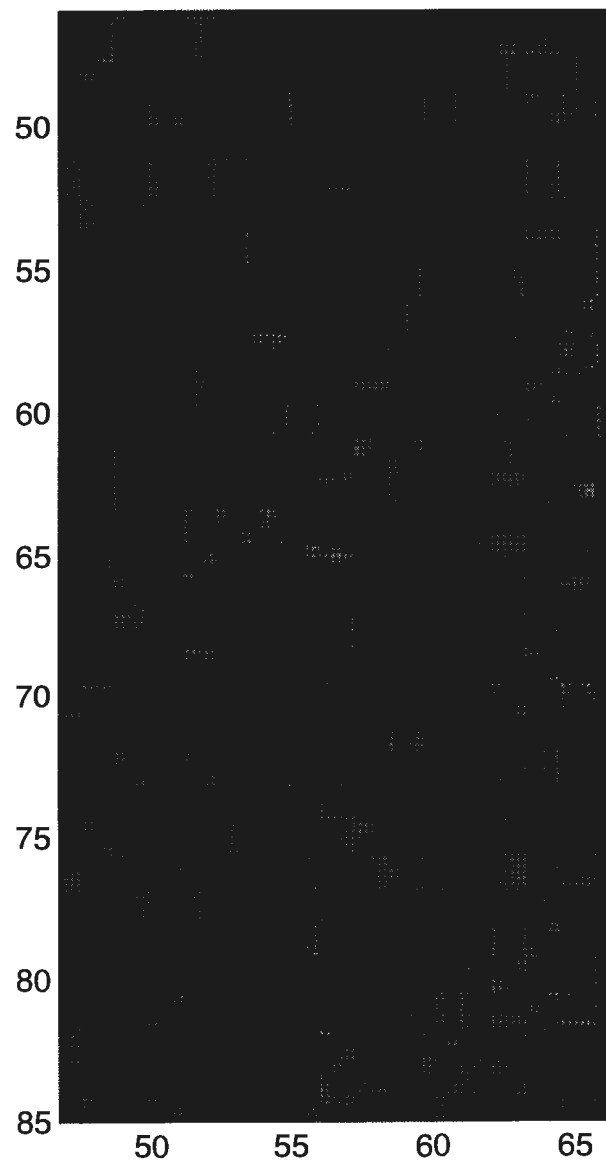


Figure 3c. JPEG Compression of Image si000206 Showing Target Location (65, 56) Magnified (50:1 Compression Image)

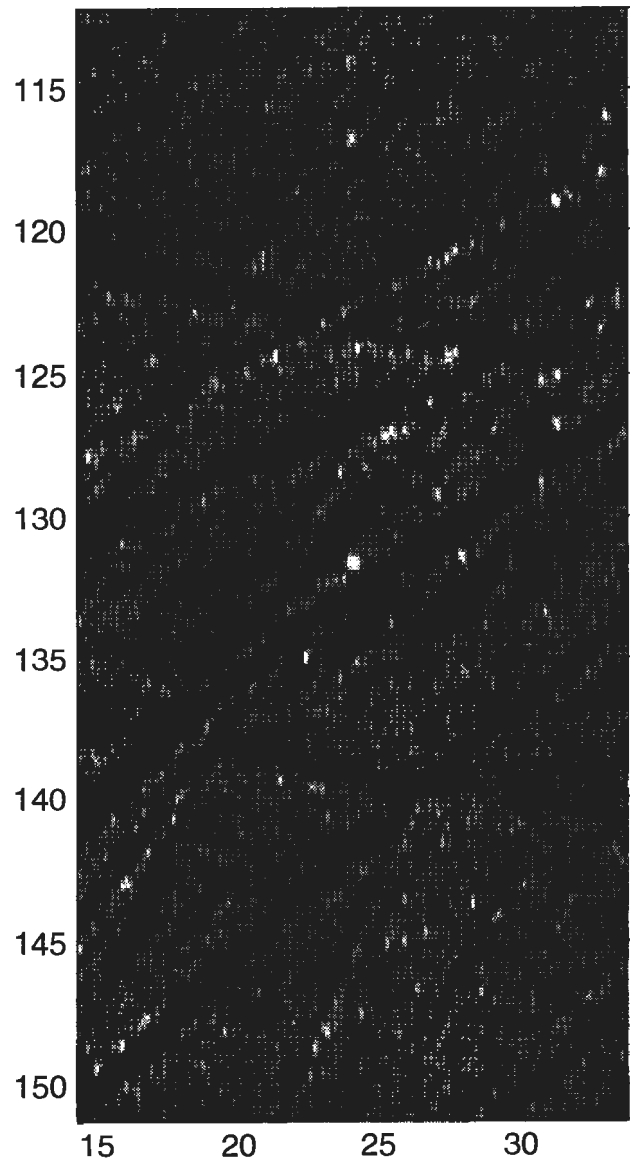


Figure 4a. JPEG Compression of Image si000206 Showing Target Location (131, 24) Magnified (Original Image)

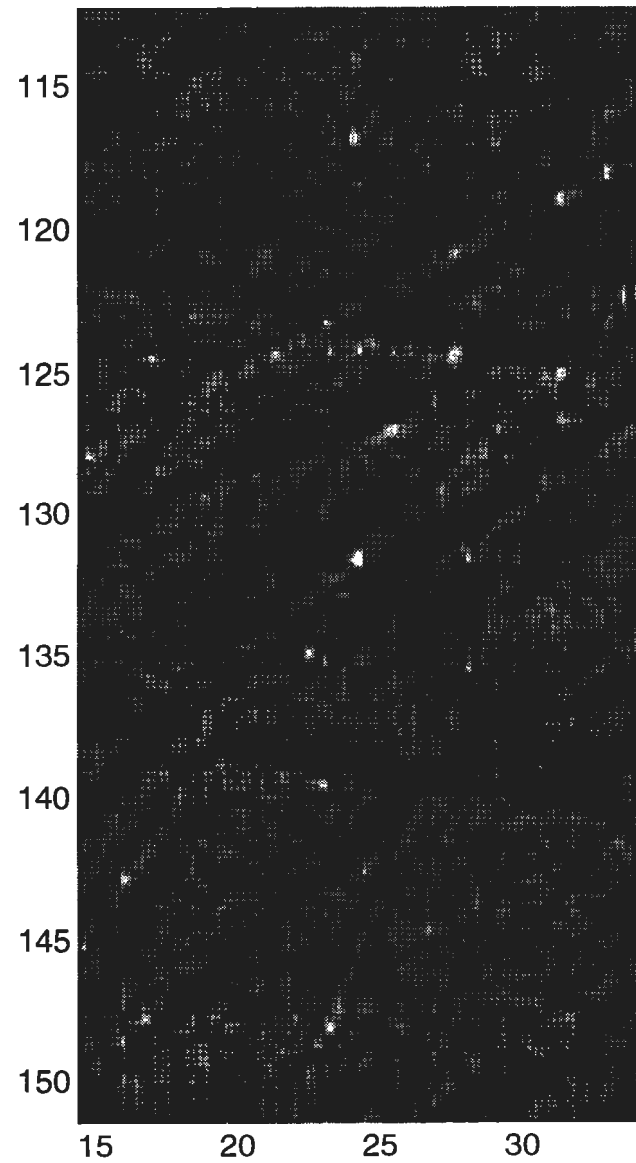


Figure 4b. JPEG Compression of Image si000206 Showing Target Location (131, 24) Magnified (225:1 Compression Ratio)

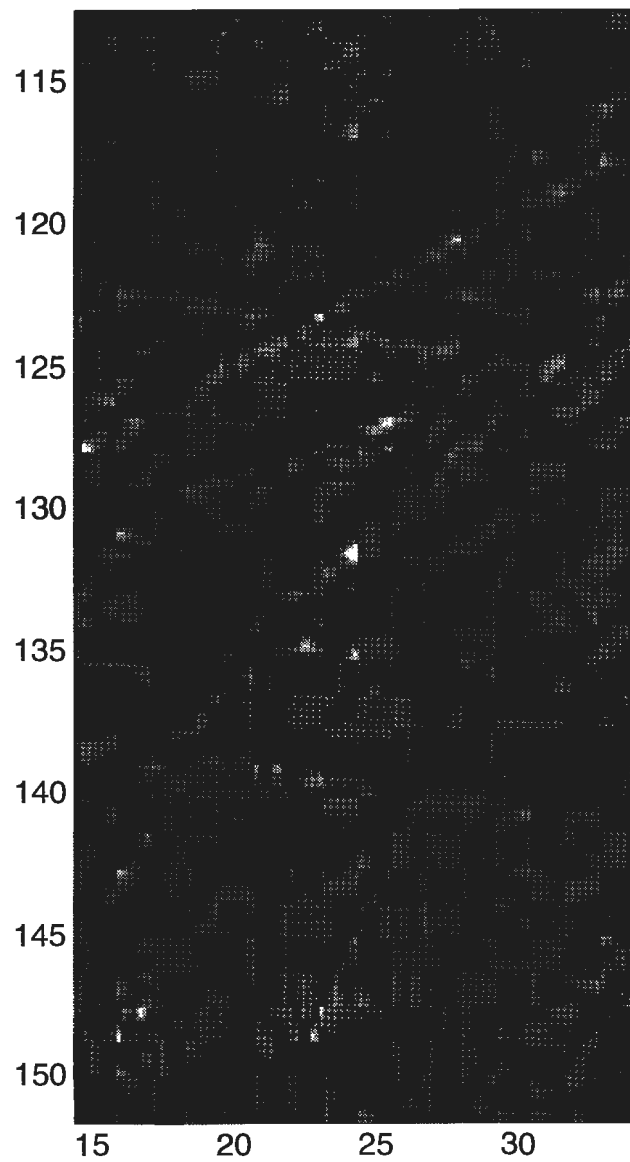


Figure 4c. JPEG Compression of Image si000206 Showing Target Location (131, 24) Magnified (50:1 Compression Ratio)

Two wavelet methods were evaluated: a biorthogonal wavelet approach and an orthogonal wavelet approach. The biorthogonal algorithm is an advanced algorithm optimizing its selection of biorthogonal basis and image processing techniques (reference 3). The orthogonal wavelet approach is a simpler approach focusing on selecting a wavelet basis that best characterized mine-like objects in side scan sonar images. The objective was to select bases and levels of processing that best preserved highlight and shadow characteristics of mine-like objects rather than using PSNR or entropy measures more suitable for an overall estimate of image fidelity.

3. ORTHOGONAL WAVELETS

PROJECTIONS ON WAVELET SPACE

Although the fleet side scan sonar image is processed in two dimensions with the orthogonal wavelet algorithm, a brief description of the theory for one dimension is presented in this section. In practice, the one-dimensional algorithm is applied twice, first to each row of the input matrix and then to each column of the row processed matrix.

The impressive success of wavelets is due mainly to the discovery of multiresolution analysis by Mallat (reference 4). Multiresolution analysis constitutes a useful functional analysis tool in wavelet theory and leads to the development of the very fast pyramid scheme to compute the wavelet coefficients. In the continuous transform, for a given function $x(t)$, the coefficients are defined as follows:

$$C_x(a, b) = \langle x, \psi_{ab} \rangle = \int_{-\infty}^{+\infty} x(t) \psi_{ab}(t) dt, \quad (1)$$

where the wavelet function $\psi_{ab}(t)$ is

$$\psi_{ab}(t) = \frac{1}{\sqrt{a}} \psi\left(\frac{t-b}{a}\right), \quad b \in R, \quad a > 0. \quad (2)$$

The parameters a and b have the effects of dilation and translation respectively. To discretize the transform in the time/frequency plane, let $a = 1/2^j$ and $b = k/2^j$, where $j, k \in Z$. The coefficients thus become

$$C_x\left(\frac{1}{2^j}, \frac{k}{2^j}\right) = 2^{j/2} \int_{-\infty}^{+\infty} x(t) \psi_{jk}(2^j t - k) dt. \quad (3)$$

Finally, discretizing in time gives

$$C_x\left(\frac{1}{2^j}, \frac{k}{2^j}\right) = 2^{j/2} \sum_{n=1}^N x[n] \psi_{jk}(2^j n - k), \quad (4)$$

where N is the length of the input vector $x[n]$. The function $\psi \in L^2(R)$ is called an orthogonal wavelet if the family $\{\psi_{jk}\}$ is an orthonormal basis of $L^2(R)$; that is,

$$\langle \psi_{jk}, \psi_{lm} \rangle = \delta_{jl} \delta_{km}, \quad (5)$$

where δ is the Kronecker delta function. Note that the wavelet function is in $L^2(\mathbb{R})$ so that it has finite support in time. This function differs from the trigonometric functions in Fourier analysis and gives the wavelet its ability to produce time information as well as frequency. One of the simplest orthogonal wavelet functions is the Haar function defined by

$$\psi(t) = \begin{cases} 1 & 0 \leq t \leq 1/2 \\ -1 & 1/2 \leq t < 1 \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

The discrete orthogonal wavelet algorithm is actually implemented as a series of convolution and decimation operations with discrete-time wavelet filter banks, such as those developed by Daubechies (reference 5). For this study, the compactly supported wavelet Daubechies 6 was adopted. The length of the wavelet chosen was fairly arbitrary. However, some of the classification work has suggested that longer wavelet filters (e.g., Daubechies 20) tend to miss some mines, and shorter wavelet filters (e.g., Daubechies 2) have many false alarms.

The discrete wavelet transform is implemented by a series of convolution and decimation operations with a pair of filters. Let $x = \{x[k]\}_{k=0}^{K-1}$ be the discrete version of input signal $x(t)$ of length $K = 2^n$, which can be either a row of the image or a column of the coefficients after the rows have been processed. In the fast discrete wavelet transform, the signal x is first decomposed into low- and high-frequency bands by the convolution-decimation (subsampling by two) operations of x with the pair of a low-pass filter $G = \{g_k\}_{k=0}^{L-1}$ and a high-pass filter $H = \{h_k\}_{k=0}^{L-1}$, where L is the length of the filter. In orthogonal wavelets, the length of the two filters is the same. The filters G and H satisfy the orthogonality conditions:

$$GH^* = HG^* = 0, \quad \text{and} \quad G^*G + H^*H = I. \quad (7)$$

G and H are called quadrature mirror filters (QMFs), which allow perfect reconstruction. The decomposition process continues iteratively on the resulting low-frequency bands and each time the high-frequency bands are left intact. The iteration stops when there is one low-frequency coefficient and one high-frequency coefficient. As a result, the frequency axis is partitioned *smoothly* and *dyadically* in an octave-band fashion, as shown in figure 5. Figure 5 shows the phase plane produced by the wavelet transform. The wavelet transform converts one-dimensional data into two-dimensional data. The horizontal, or t , axis can be labeled by time or position, depending on the nature of the data, and increases to the right. The vertical, or f , axis is usually labeled frequency, or scale, and increases upward. Different spatial resolutions are given to different frequency bands. Low frequency with low spatial resolution is at the bottom, while, towards the top of the figure, the frequency resolution is decreased and the spatial resolution is increased. The entire phase plane is covered by disjoint cells of equal area, which are called Heisenberg cells (reference 6). The uncertainty principle can be interpreted as a rectangular cell located around a position in the phase plane, (t, f) , that represents an uncertainty region associated with (t, f) . The total number of cells is equal to the dimension of the input vector. Each cell can be shaded in proportion to the magnitude of the corresponding wavelet coefficient to create a gray-scale image of the phase plane. It is clear that this type of gray-scale quantization procedure of cells conforms with the uncertainty principle.

WAVELET PACKET

For the wavelet packet transform (reference 7), the high-frequency band, which is left intact during each iteration of the wavelet transform, is also decomposed into finer frequency bands. Figure 6 depicts the entire wavelet packet decomposition tree. The first level represents the original signal x . The second-level decomposition generates Hx , labeled x_h , and Gx , labeled x_g . Gx represents the low-frequency band, and Hx represents the high-frequency band. Applying the low-pass filter G and the high-pass filter H to both low-and high-frequency bands obtains four frequency bands H^2x , GHx , HGx , and G^2x ordered in decreasing frequency, which is the third level of decomposition. This decomposition process continues to the maximum depth of $J + 1$, where $J = \log_2(n)$, and finer frequency resolutions are obtained toward lower levels. Because each decomposition level generates equal boxes corresponding to uniform partition of the frequency axis, the time axis is also windowed uniformly. Hence, the extent of the support of each basis function is essentially constant for a decomposition level. Note that each level forms an orthonormal basis on which the side scan sonar image is intended to be projected.

The full wavelet packet transform produces a wavelet packet tree structure containing many more coefficients than are needed to reconstruct the image. A basis vector of coefficients should be selected. The basis vector of transform coefficients can be constructed by selecting all the coefficients in a level, referred to as “a level basis,” or by selecting coefficients from different levels to obtain a “best” basis for the input data vector x . Best can be in whatever measure desired. Certain rules that constrain the selection of coefficient sets having parent-child relationships within the wavelet packet decomposition tree apply (reference 7). For the fleet side scan sonar application, a level basis that best characterized the spatial and spectral characteristics of the targets was selected. The transform vector contains the same number of coefficients as the dimension of the input vectors from the original image, and does not lead to compression. However, the forward transform increases the amount of energy contained in some individual coefficients and decreases others. Increasing the magnitude of individual coefficients increases the data requirements of the transform, but, because the energy is conserved in the transform, most of the remaining coefficients are very small, approaching zero. Because deleting small coefficients does not significantly affect the total energy in the image, it will not cause significant distortion in the reconstructed image. The surviving coefficients now comprise the transform vector and can be encoded for compression using a zero run-length encoding technique.

The bases chosen were level three for both the rows and columns of the image. Using these bases, each coefficient can be thought of as representing an eight by eight square of pixels in the image. Part of the reasoning for using these bases is that eight by eight pixels match well with the typical size of mines in the images. Figure 7a shows how a typical mine might be partitioned by coefficients at level 3. At this level of resolution, each coefficient covers a good portion of the mine along with the frequency components associated with its edges but is not lost within the interior of the mine. Figure 7b shows the resolution of coefficients at level 2. At this level, there are many more coefficients that lie on the mine; however, some of them are in the interior of the mine and will not capture the frequency components because of the edges. Also at level 2, there are many more coefficients to be retained to preserve the image of the mine. Figure 7c shows the resolution of coefficients at level 4. Here there are fewer coefficients (but more frequencies) that lie on the mine. While fewer must be retained, these coefficients contain a significant amount of information from outside the mine. This occurrence can cause the coefficients to be smaller and less likely to be retained in the compression algorithm and also tends to blur the mine upon reconstruction of the image.

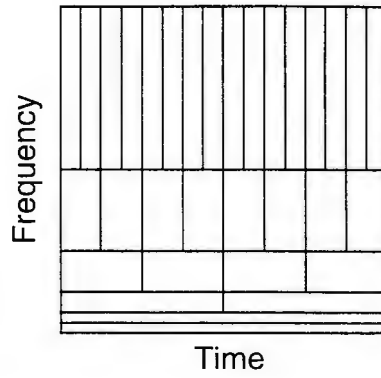


Figure 5. Two-Dimensional Wavelet Spatial/Temporal Representation of a One-Dimensional, 32-Point Signal

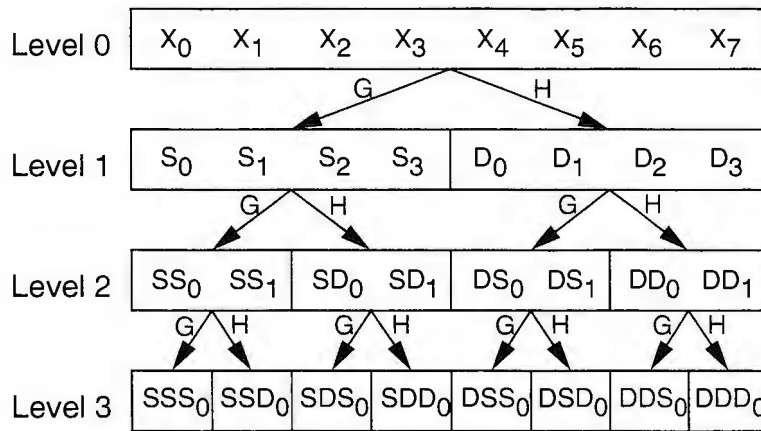


Figure 6. Wavelet Packet Transform Coefficient Tree

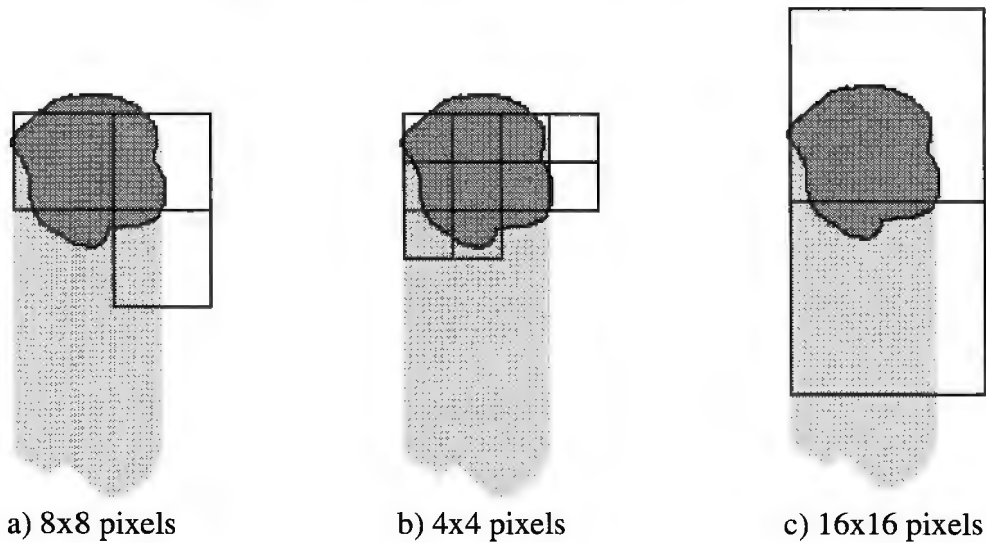


Figure 7. Image Partitioning Using (a) Coefficients from Level 3, (b) Coefficients from Level 2, and (c) Coefficients from Level 4

ORTHOGONAL TRANSFORM ALGORITHM

The algorithm for the orthogonal wavelet case is as follows: the one-dimensional wavelet packet algorithm is applied twice, first to each row of the input matrix and then to each column of the row processed matrix. Only the level 3 basis coefficients for the lower half of the frequency spectrum are retained, which provides some filtering of high-frequency noise. Next, the coefficients are sorted by magnitude for the spectrum of interest. Processing time is substantially reduced by only processing the pyramid algorithm to level 3 and limiting the spectrum of interest to the lower half of the image bandwidth. The number of coefficients required to give the desired compression ratio is calculated, and only this number of the largest coefficients, along with their locations in the coefficient matrix, is retained in the compressed image file. Zero run-length encoding is currently used to encode coefficient locations. For reconstruction, the surviving coefficients are placed into their proper locations in the coefficient matrix while the rest of the coefficients are set to zero. The inverse wavelet packet transform is performed on each column, replacing data that were there, and then it is performed on each row to produce the reconstructed image that is processed by the classification algorithm.

OVERALL PERFORMANCE

The fleet side scan sonar images were compressed 25:1, 50:1, and 100:1 using the orthogonal wavelet technique. The original and reconstructed images for each of the compression ratios were generated for this example, file si000206, and are shown in figure 8. The regions containing each of the two targets are magnified and shown in figures 9 and 10. At a compression ratio of 25:1, the most noticeable difference from the original image is the reduction of high-frequency texture information from the background. This effect increases at 50:1 and at 100:1. The texture information has been mostly removed from the near-range region. A significant weakening of the highlight or shadow contrast is not observed in either of the targets in figures 9 and 10 because the compression ratio is increased from 25:1 to 100:1.

4. BIORTHOGONAL WAVELETS

BIORTHOGONAL TRANSFORM

The advantage of the wavelet transform is that its complexity is of order n ($O(n)$), where n is the number of points in the data, and that the implementation is through finite impulse response (FIR) filters. The primary differences among the various types of wavelet transforms are the length of the filters, the required precision of the filter coefficients, and the relationship between the forward and inverse filters. Shorter filters result in fast computation. Integer filter coefficients result in smaller computational errors and reduced calculation complexity. The relationship between the forward and inverse filters can provide a means of minimizing the distortion introduced in the lossy compression process and of reducing computational complexity.

All wavelet transforms provide perfect reconstruction. Orthogonal wavelet transform filters implement the same filtering at the forward and inverse transforms. Biorthogonal wavelet transforms implement different filters at the forward and inverse transforms. The advantage of the biorthogonal transform is its “pre-emphasis” and “de-emphasis” operations. On the forward end, the transform enhances certain image information and thus provides a method for better extraction of important detail from typical image data. On the inverse end, the transform smoothes the image information and thus reduces the impact of lost information on the quality of the reconstructed image.

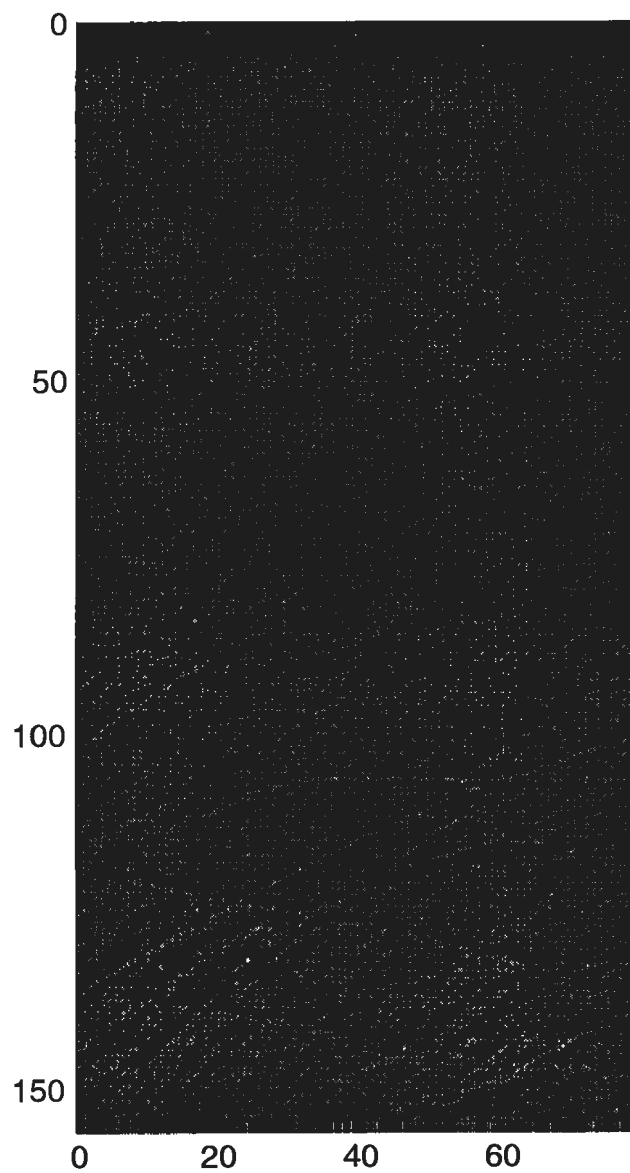


Figure 8a. Orthogonal Wavelet Compression of Image si000206 (Original Image)

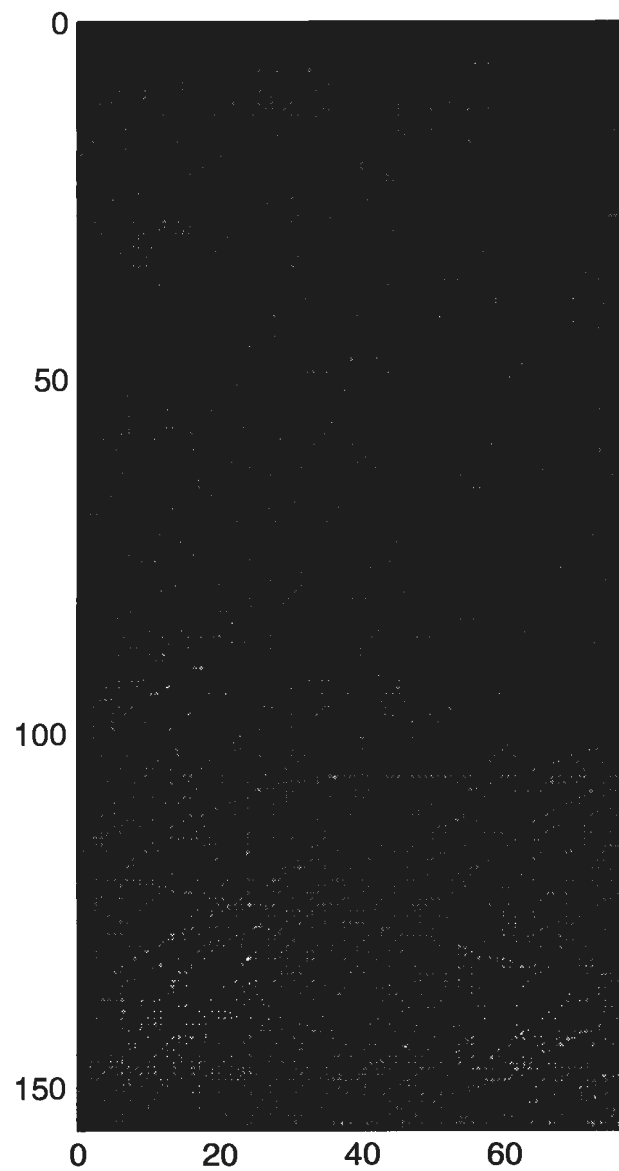


Figure 8b. Orthogonal Wavelet Compression of Image si000206 (25:1 Compression Ratio)

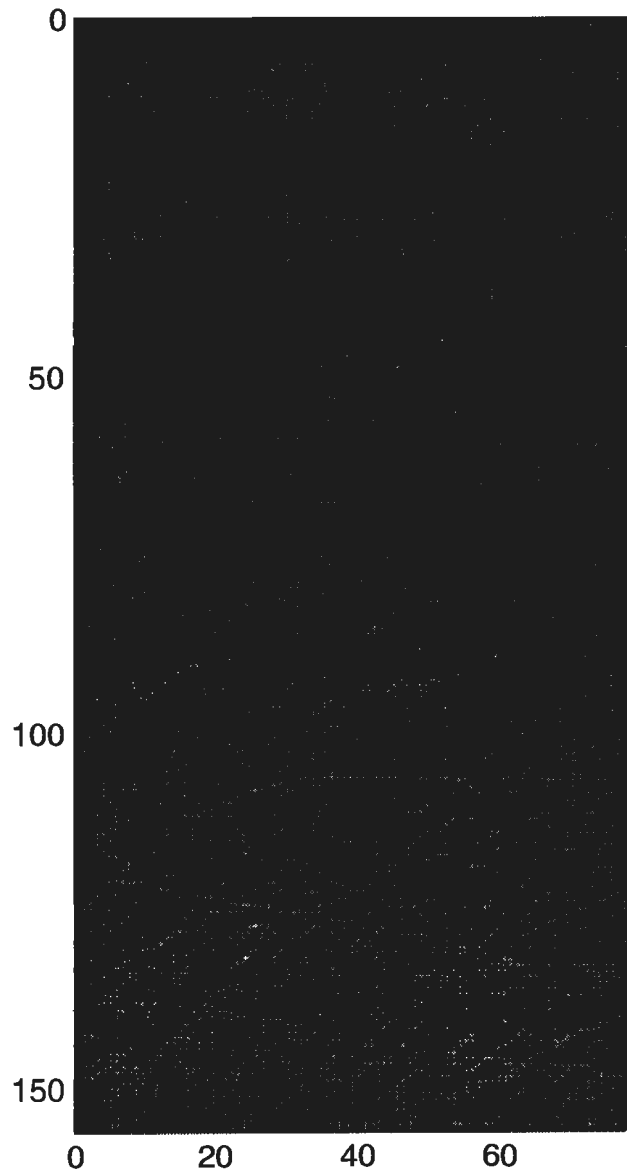


Figure 8c. Orthogonal Wavelet Compression of Image si000206 (50:1 Compression Ratio)

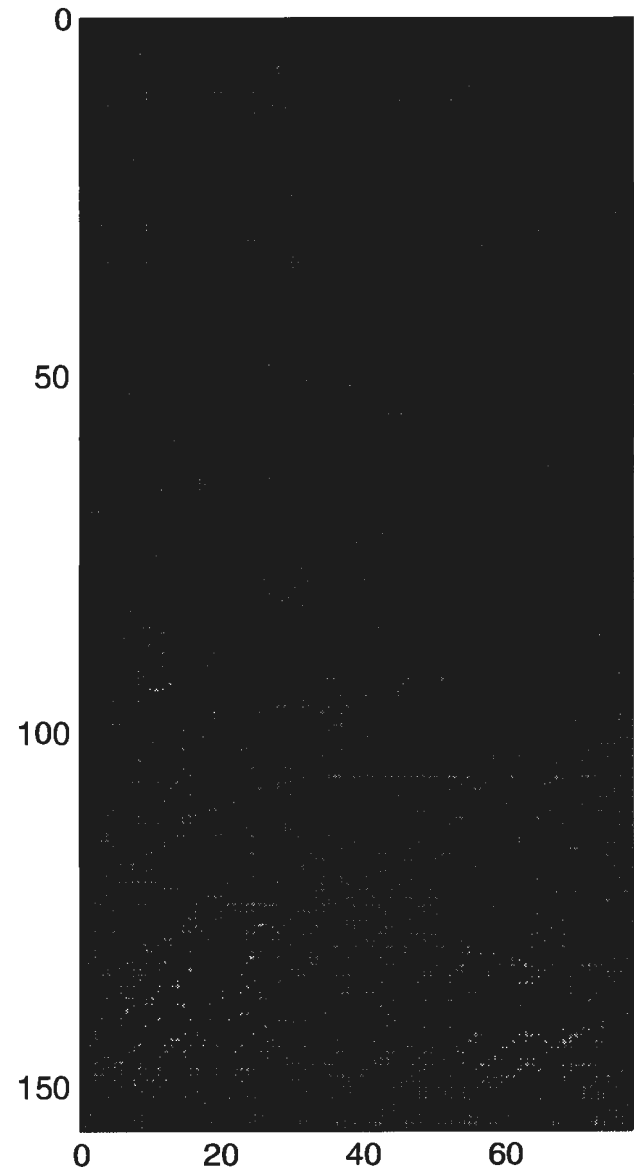


Figure 8d. Orthogonal Wavelet Compression of Image si000206 (100:1 Compression Ratio)

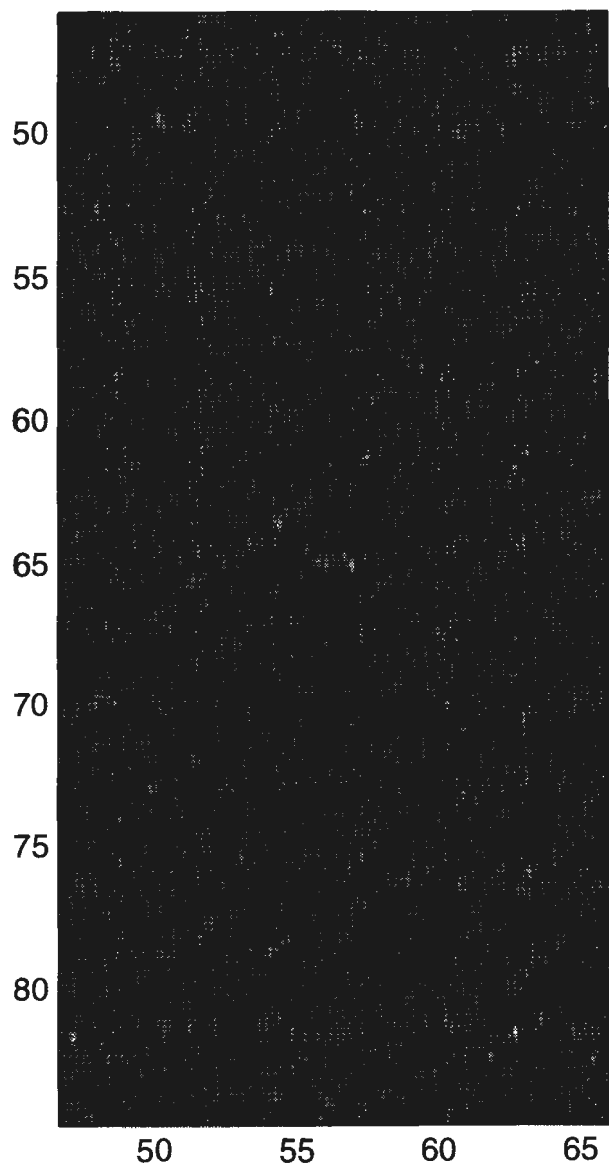


Figure 9a. Orthogonal Wavelet Compression of Image si000206 Showing Target Location (65, 56) Magnified (Original Image)

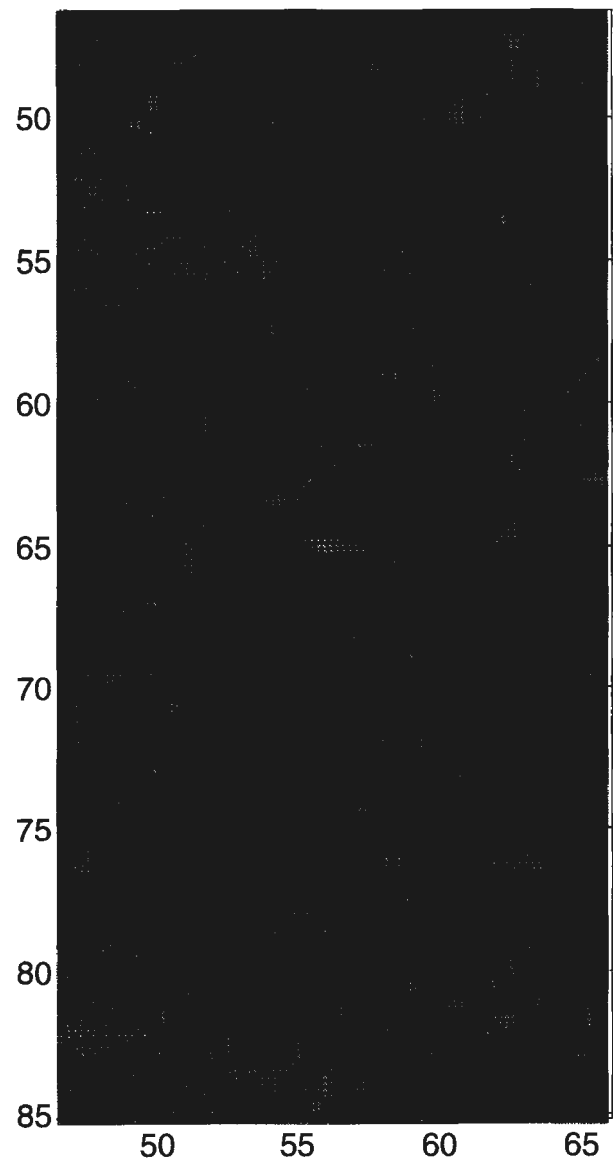


Figure 9b. Orthogonal Wavelet Compression of Image si000206 Showing Target Location (65, 56) Magnified (25:1 Compression Ratio)

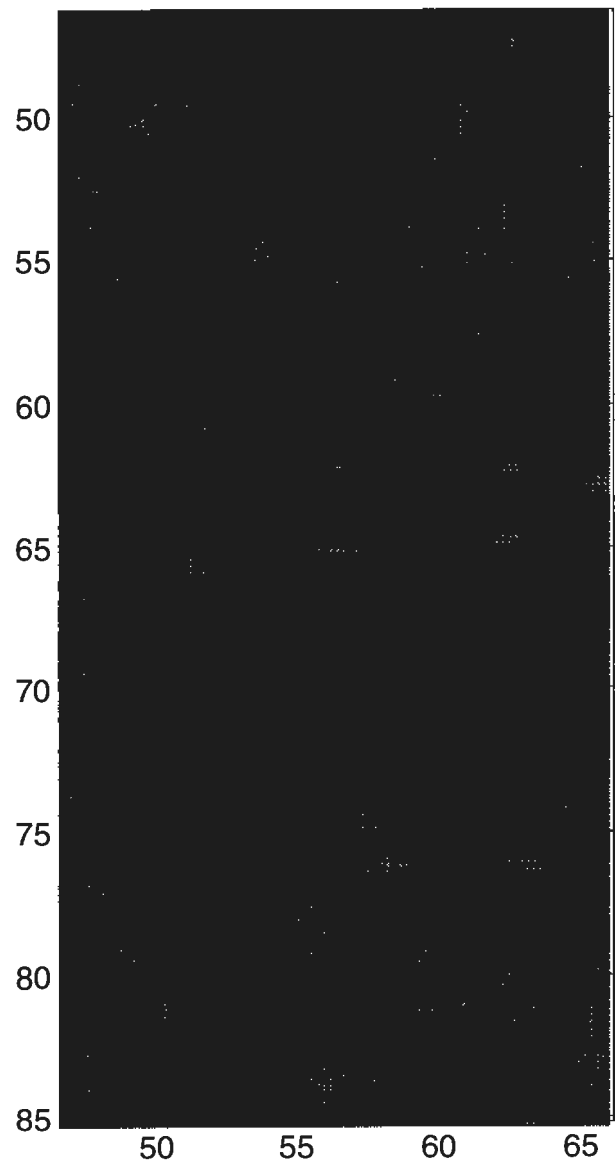


Figure 9c. Orthogonal Wavelet Compression of Image si000206 Showing Target Location (65, 56) Magnified (50:1 Compression Ratio)

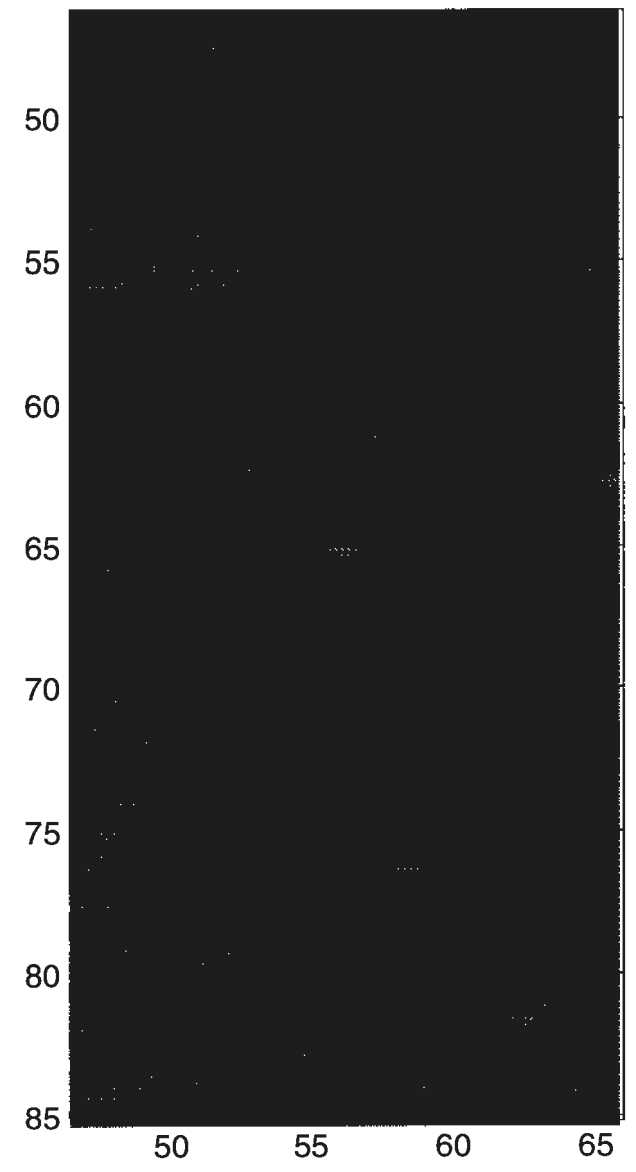


Figure 9d. Orthogonal Wavelet Compression of Image si000206 Showing Target Location (65, 56) Magnified (100:1 Compression Ratio)

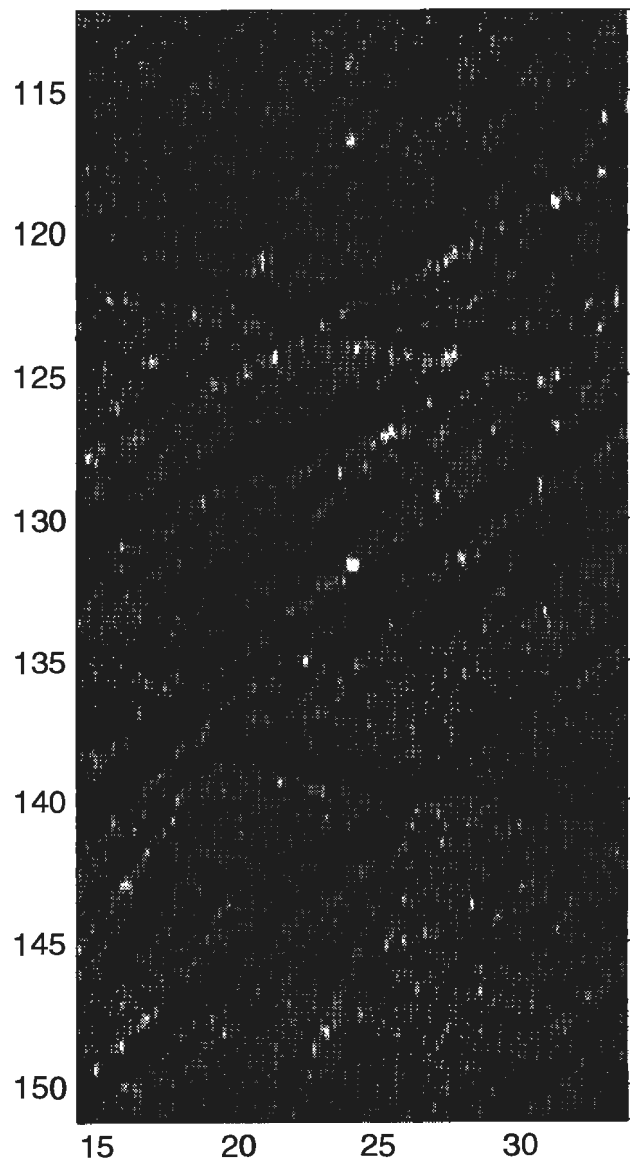


Figure 10a. Orthogonal Wavelet Compression of Image si000206 Showing Target Location (134, 24) Magnified(Original Image)

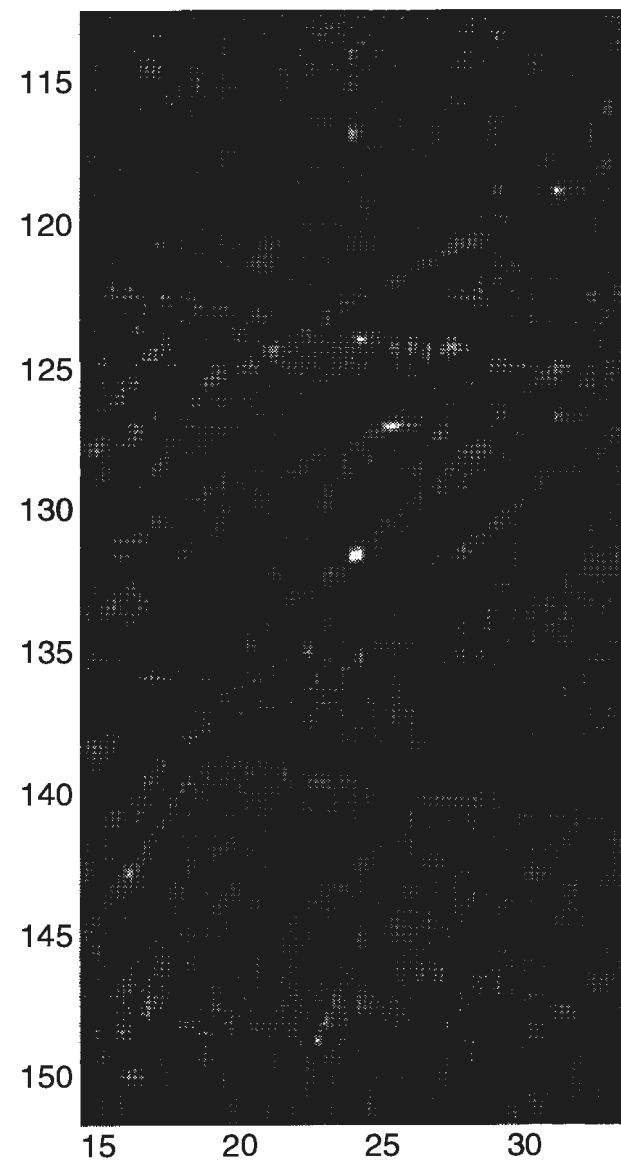


Figure 10b. Orthogonal Wavelet Compression of Image si000206 Showing Target Location (134, 24) Magnified (25:1 Compression Ratio)

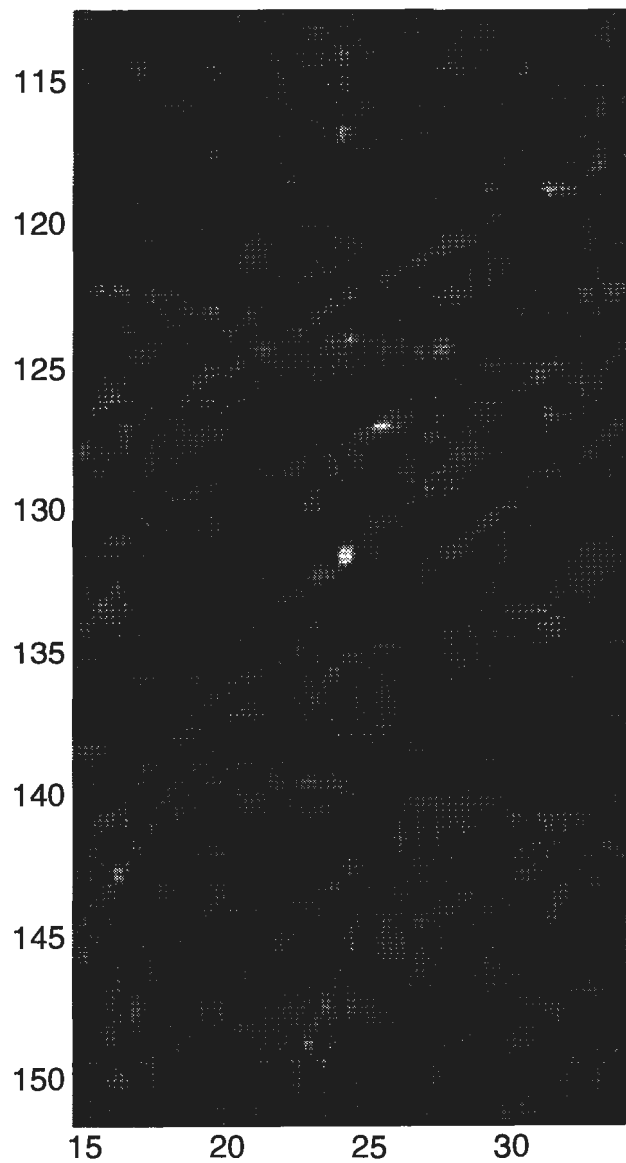


Figure 10c. Orthogonal Wavelet Compression of Image si000206 Showing Target Location (134, 24) Magnified(50:1 Compression Ratio)

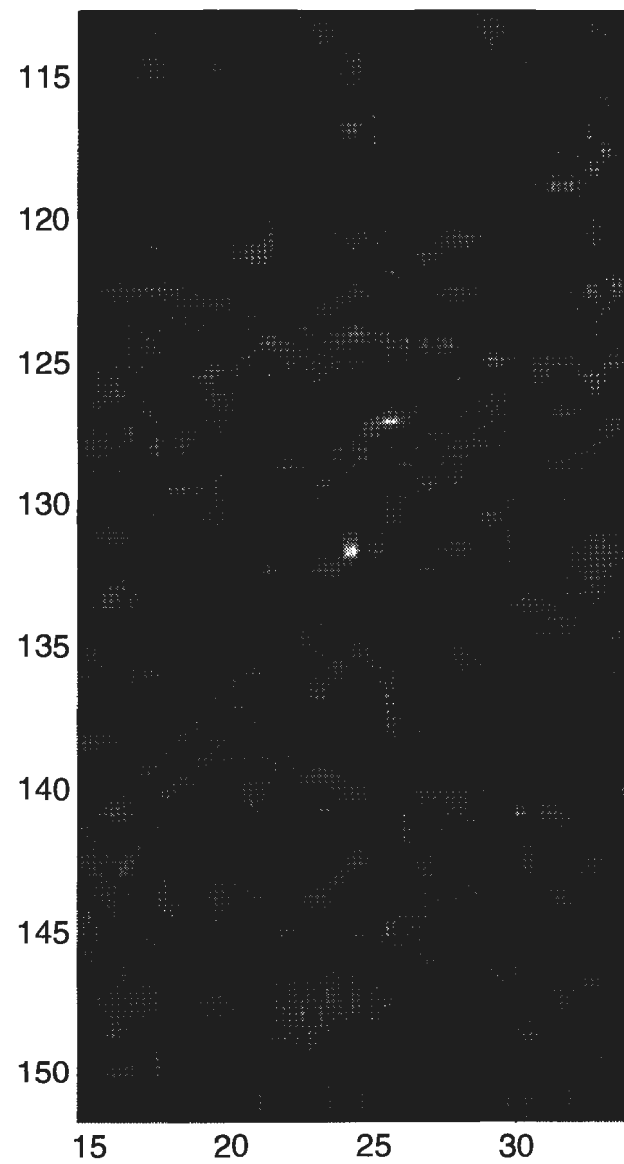


Figure 10d. Orthogonal Wavelet Compression of Image si000206 Showing Target Location (134, 24) Magnified (100:1 Compression Ratio)

QUANTIZATION

There are several approaches to quantization. Quantization methods may be uniform, nonuniform, or adaptive, depending on the transform level and the statistics of the image. If the image data set is small and the general image statistics are well understood, then the quantization can be predetermined.

ENCODING

The encoding of transformed and quantized image data is achieved through some lossless encoding process consisting of two basic operations:

1. encoding the addresses of the surviving coefficients and
2. encoding the values of the surviving coefficients.

The most popular approach to address coding is run-length encoding (reference 8); the most popular approach to coefficient value coding is Huffman encoding (reference 9). Further compression may be achieved by Huffman encoding of the run-lengths. Other methods, including arithmetic encoding (reference 10) and zerotree encoding (reference 11), are also available.

ENHANCEMENTS

Several enhancements exist and may be performed with low additional complexity on the transform coefficients themselves during compression or decompression. These enhancements can add the following capabilities:

1. progressive transmission
2. region-based quantization
3. contrast equalization
4. edge enhancements
5. magnification
6. sharpening
7. smoothing.

This approach saves at least the cost of rescanning the image data for separate processing. In addition, the number of operations can often be reduced further by performing the necessary calculations on the surviving coefficients only.

Other enhancements (not listed) may be matched more closely to the needs of the automatic detection/classification algorithms. For example, there are several methods to reduce background clutter noise through simple operations on the transform coefficients. Additionally, preliminary feature detection can be performed in the wavelet domain at a modest additional complexity cost. The identified features may be less severely quantized, and hence better preserve potentially important features. Also, the location of the features could be provided as side information to the classifier.

ARCHITECTURE

The general construction of the biorthogonal transform lends itself to a parallel processing environment. Another advantage of this architecture is that the image may be processed in

sections with little loss in performance, thus allowing the processing of very large images in limited memory environments. The compact nature of the basic filter kernel allows efficient low memory/low latency architectures.

OVERALL PERFORMANCE

The biorthogonal wavelet-based compression algorithm is an extremely fast and efficient compression engine developed for general image compression. The flexibility and architecture of the algorithm are easily adapted to application-specific requirements. Current performance benchmarks for gray-scale images are as follows (for an independent evaluation of performance please see <http://www.summus.com> or reference 12):

Quality-1.5 to 2 times better than JPEG at moderate compression ratios based on PSNR measures, and significantly better at high compression ratios;

Compression Ratios-up to 300:1 for data useful to the human observer;

Speed-less than 1 second to compress a 512-by 512-pixel image on a 486-66 MHz processor, and speed scales linearly with processor/memory speed.

Further performance improvements scale with the power of the processor and with new developments in image compression technology. The algorithm also provides several enhancement techniques for efficient improvement of data clarity. Fleet side scan sonar images were compressed by 25:1, 50:1, and 100:1 using the advanced biorthogonal wavelet technique, a high-performance technique applicable to most image compression problems. Using the biorthogonal algorithm, the original and reconstructed images for each of the compression ratios were generated for file si000206 and are shown in figure 11. The regions containing each of the two targets are magnified and are shown in figures 12 and 13. These figures show the significant improvement in fidelity when compared to the JPEG results shown in figures 2 through 4.

5. AUTOMATED DETECTION AND CLASSIFICATION

The “goodness” of the compression methods could have been evaluated subjectively by an expert sonar operator by evaluating operator ability to detect and classify mines while minimizing the number of false target calls. However, to be more objective, an automated detection and classification (D/C) algorithm was used where the probability of detection and classification (PdPc) as well as the average number of false alarms per image (FAI) were used as figures of merit for performance. If the PdPc/FAI for the compressed images is similar to that for the original, one can reasonably assume that pertinent features are being preserved by the compression methods.

Classification performance varies greatly as a function of sonar type, environment, and mine type. For the last seven years, NSWCDD, CSS, has used sonar images from three different side-looking sonars to evaluate D/C algorithms developed in-house and by many others in industry and universities. These databases were selected because they demonstrate a variety of problematic issues related to environment, especially bottom texture and sonar motion. Mine threats in these databases are bottom mines, which provide a significantly greater challenge to detect and classify than volume mines. Performance results referred to in this TM are for single pass (one look), which is of interest because of the implied higher search rate. Also the performance of an expert sonar operator was available for these databases, thus providing

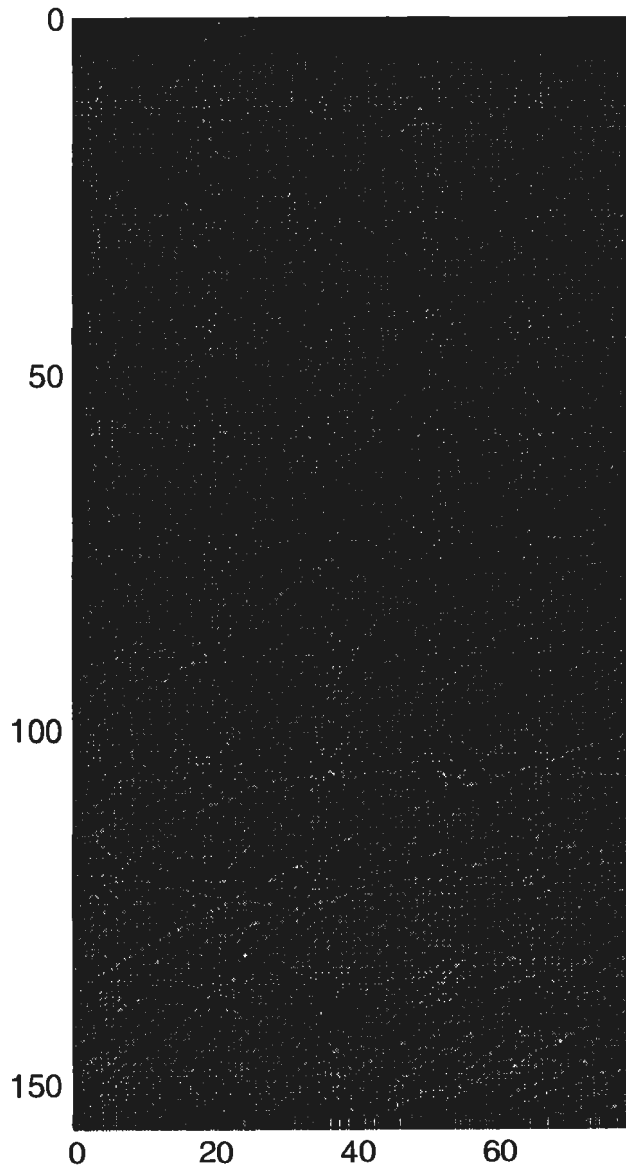


Figure 11a. Biorthogonal Wavelet Compression of Image si000206 (Original Image)

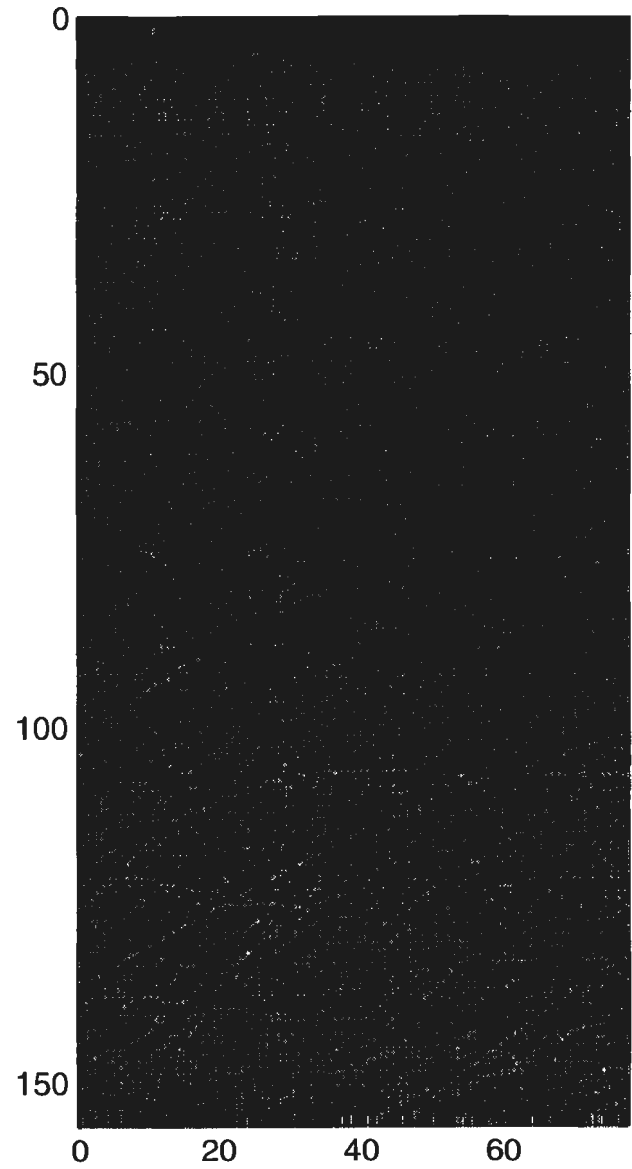


Figure 11b. Biorthogonal Wavelet Compression of Image si000206 (25:1 Compression Ratio)

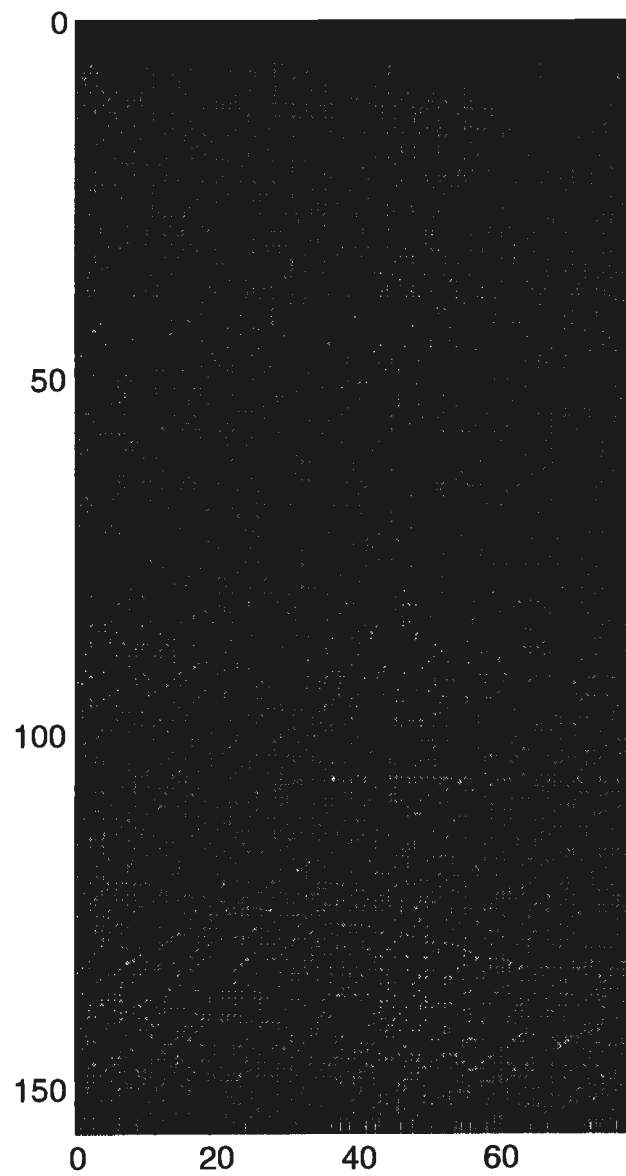


Figure 11c. Biorthogonal Wavelet Compression of Image si000206 (50:1 Compression Ratio)

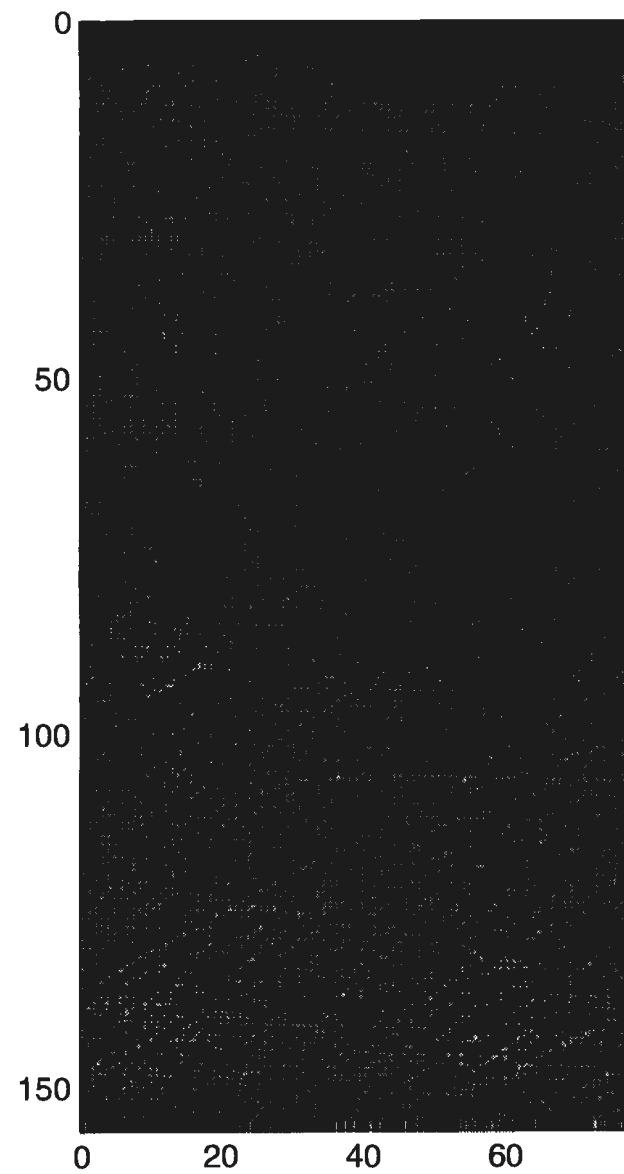


Figure 11d. Biorthogonal Wavelet Compression of Image si000206 (100:1 Compression Ratio)

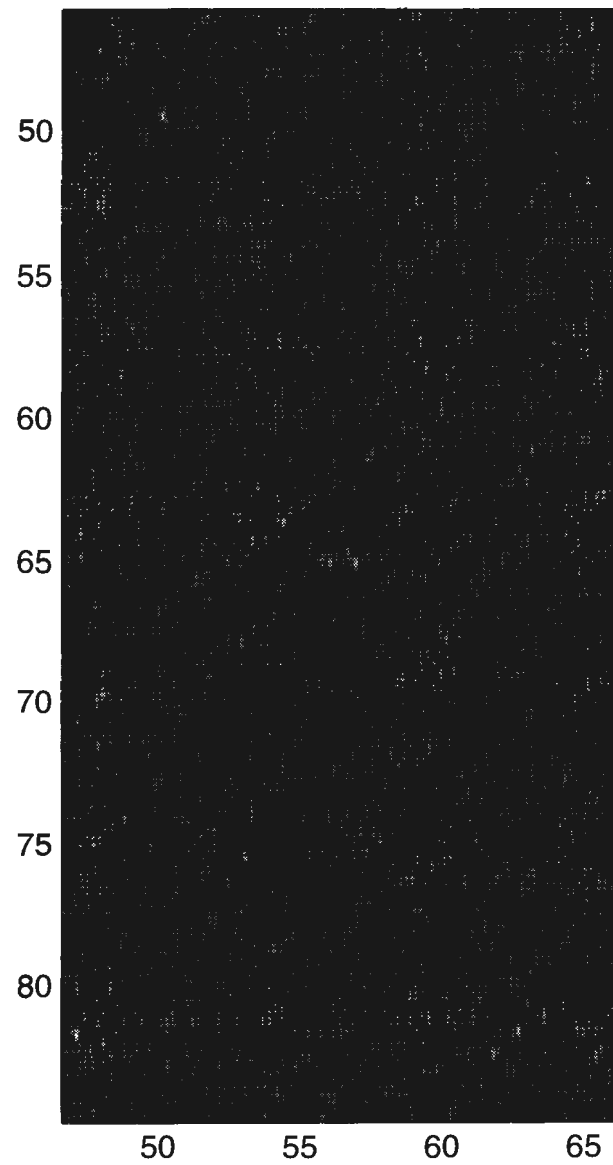


Figure 12a. Biorthogonal Wavelet Compression of Image si000206 Showing Target Location (65, 56) (Original Image)

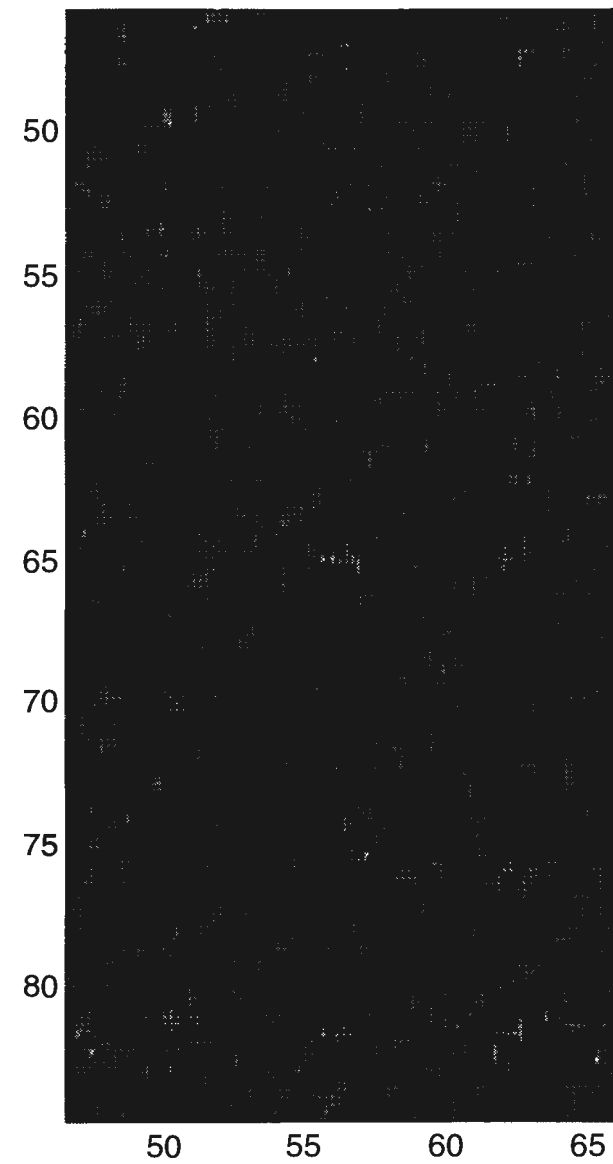


Figure 12b. Biorthogonal Wavelet Compression of Image si000206 Showing Target Location (65, 56) (25:1 Compression Ratio)

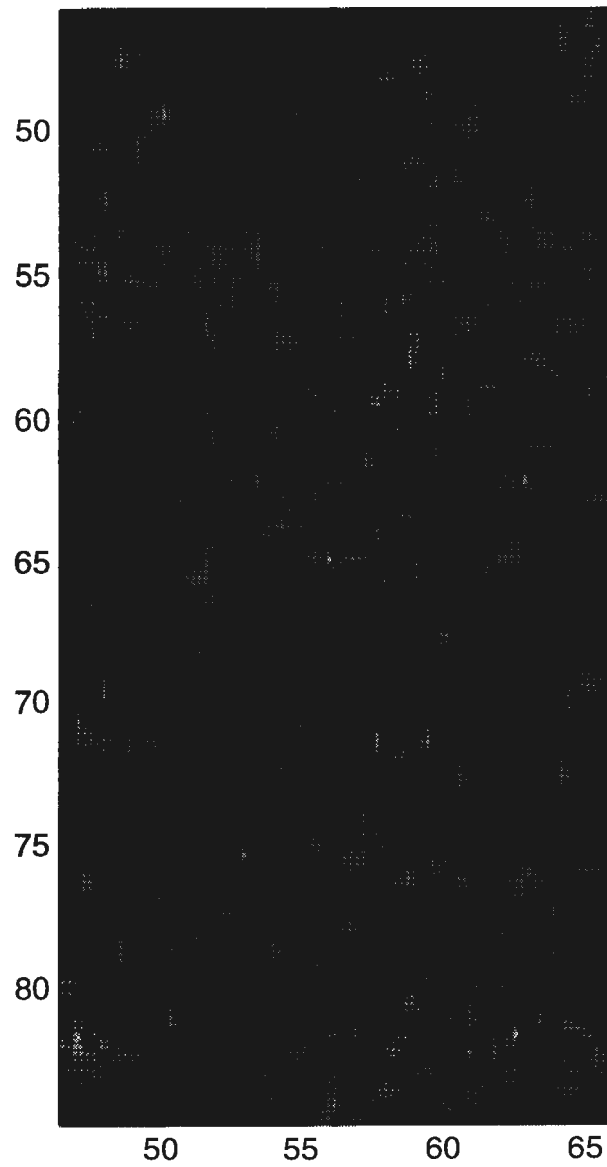


Figure 12c. Biorthogonal Wavelet Compression of Image si000206 Showing Target Location (65, 56) Magnified (50:1 Compression Ratio)

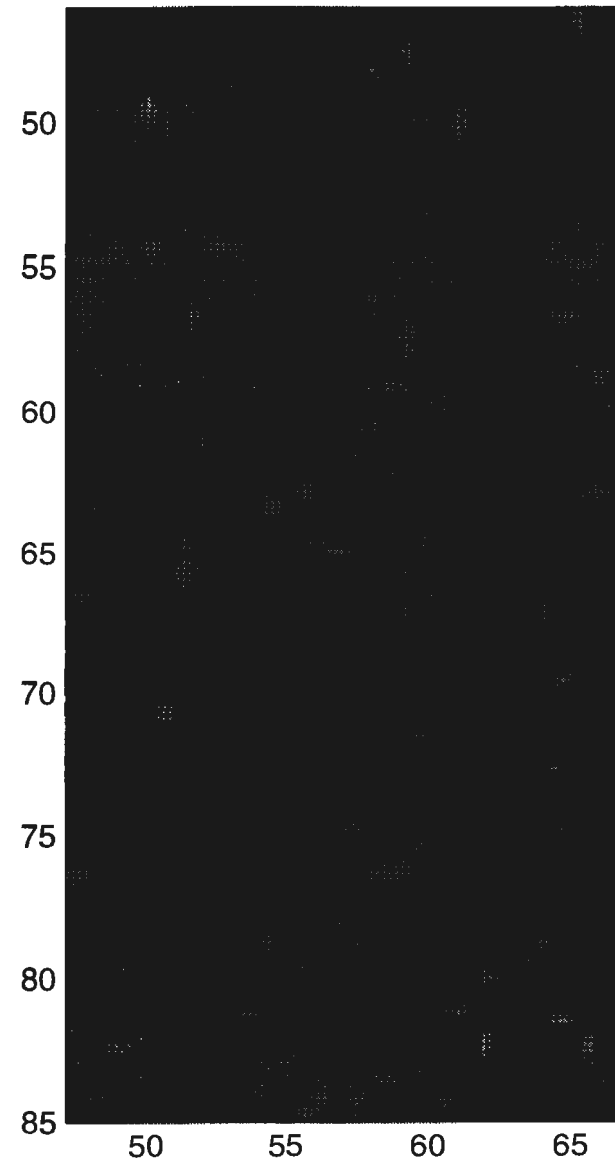


Figure 12d. Biorthogonal Wavelet Compression of Image si000206 Showing Target Location (65, 56) Magnified (100:1 Compression Ratio)

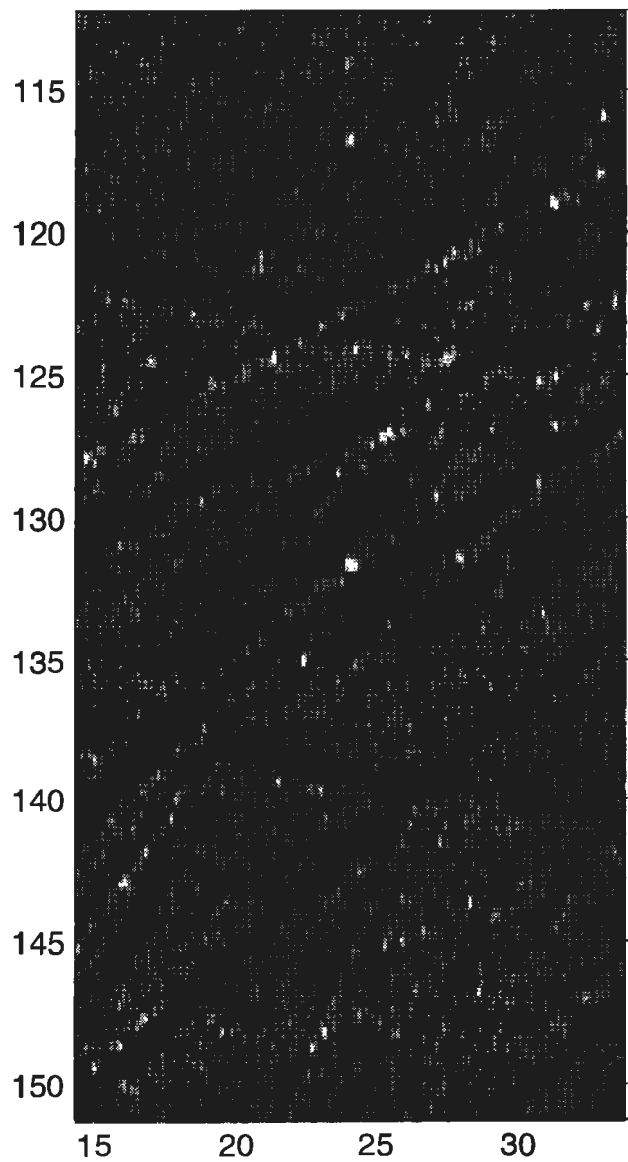


Figure 13a. Biorthogonal Wavelet Compression of Image si000206 Showing Target Location (131, 24) Magnified (Original Image)

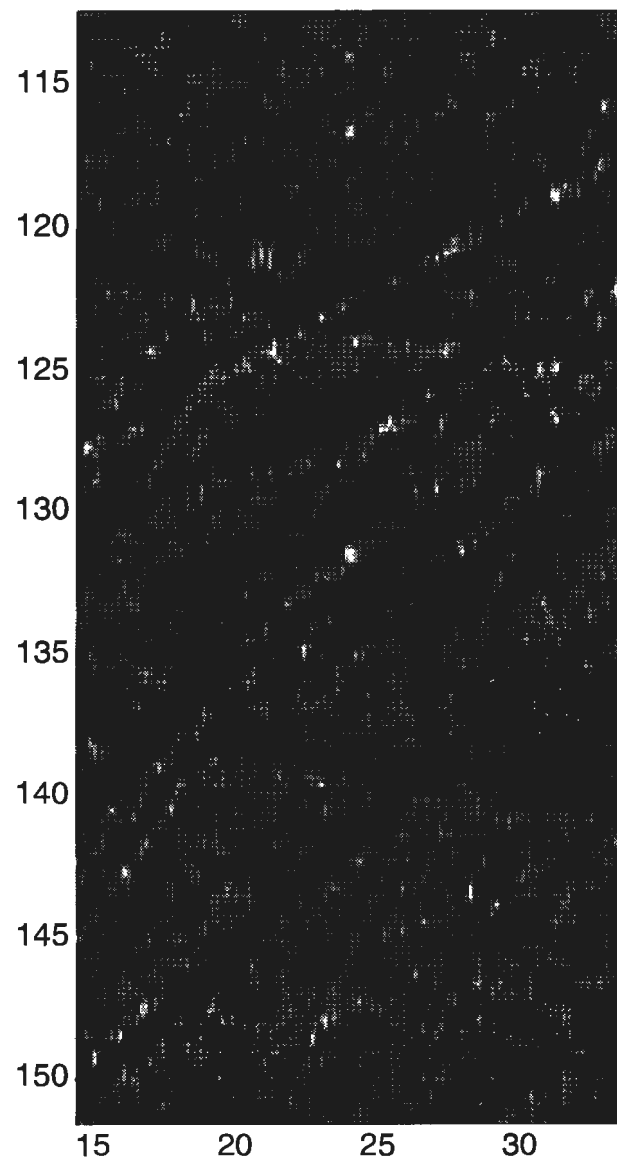


Figure 13b. Biorthogonal Wavelet Compression of Image si000206 Showing Target Location (131, 24) Magnified (25:1 Compression Ratio)

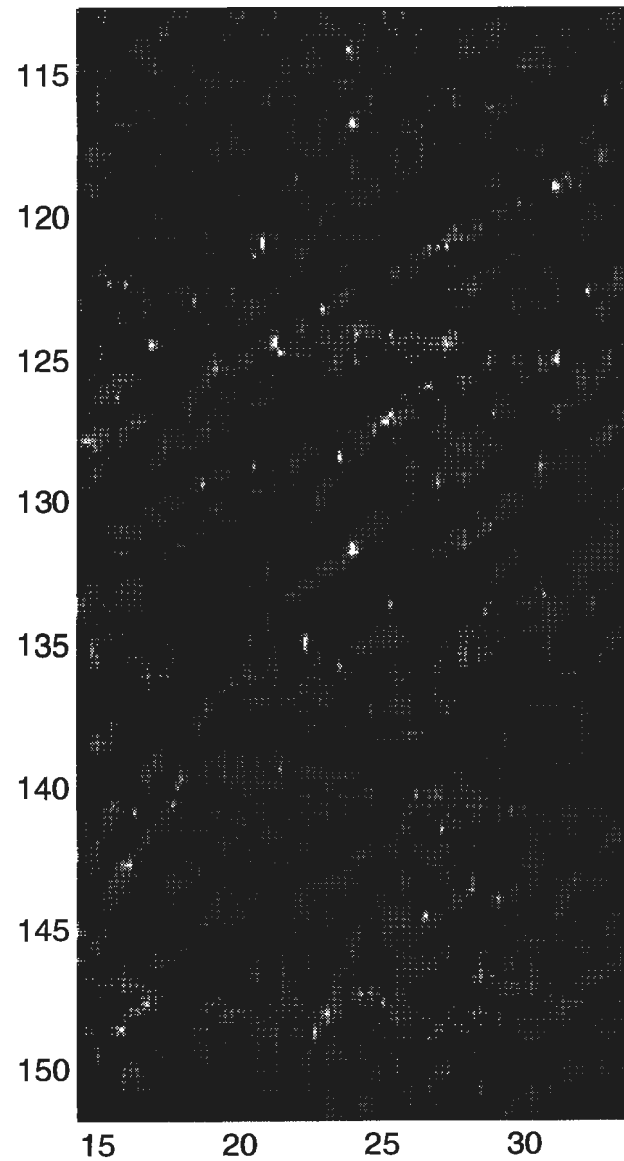


Figure 13c. Biorthogonal Wavelet Compression of Image si000206 Showing Target Location (131, 24) Magnified (50:1 Compression Image)

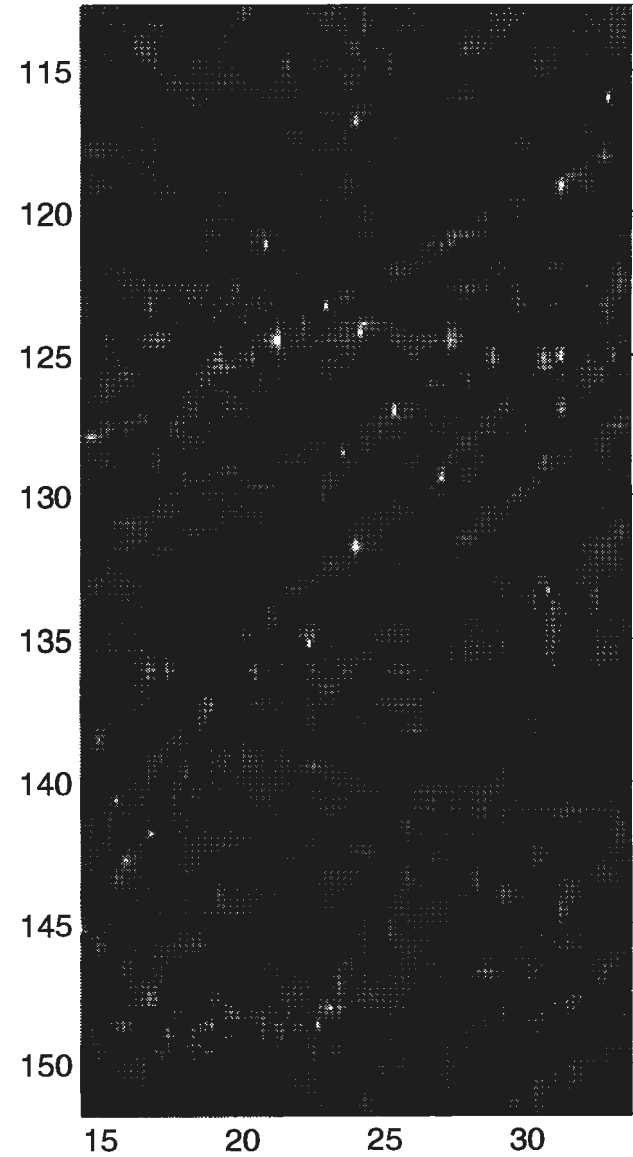


Figure 13d. Biorthogonal Wavelet Compression of Image si000206 Showing Target Location (131, 24) Magnified (100:1 Compression Ratio)

guidance to what kind of performance might be achievable by an automated algorithm. Over the last 4 years of research, the CSS D/C algorithm, described below, has produced the same PdPc for two of the three databases as the expert sonar operator at statistically equivalent FARs. For the third database, a 10-percent increase in PdPc was realized at the same FAR. Of course these three databases do not represent all interesting scenarios, so one must be cautious in generalizing any conclusion. To date, utilization of small databases seems the most cost-effective means to evaluate and compare algorithms because of the prohibitive cost of collecting and processing large databases. Also, simulations are expensive and do not capture all the critical dynamics.

The fleet side scan sonar database was used in evaluating the biorthogonal and orthogonal compression methods. The database consists of 60 images containing 33 mine signatures and many false target signatures. From this original data set, six additional data sets of reconstructed images were generated from images compressed with the biorthogonal and orthogonal wavelet methods at compression ratios of 25:1, 50:1, and 100:1.

Each of the seven data sets (1 original and 6 reconstructed) was evaluated by the D/C algorithm shown in figure 14. The algorithm was developed recently at CSS and consists of four stages: (1) detection, (2) feature extraction, (3) optimal feature selection, and (4) classification. The two classifiers used in the algorithm are the K-nearest neighbor attractor-based neural network (KNN) and the optimal discrimination filter classifier (ODFC). This algorithm was used to determine if the two compression methods preserve critical classification features in the application of finding sea mines in side scan sonar imagery. A brief synopsis of each of the four stages of the algorithm is given below and detailed discussion is found in reference 13.

DETECTION

The purpose of the detection stage is to scan the entire image and reduce the data to a small number of candidate mine-like regions. In this study, the detection stage produced an average of from 1.4 false alarms per image for the no-compression case to 8.1 false alarms per image for the 100:1 compression case at a probability of detection (Pd) of 91 percent. This part of the algorithm is the most computationally intensive because the entire image is scanned. However, on a per-pixel basis, the computation requirement is very low. Its goal is to reduce the number of mine-size regions that have to be processed by the subsequent classification stages that are computationally much more intensive on a per-pixel basis than the detection stage. This goal is reasonable because the classification stages are extracting much more information about the targets. Typically 20 detections or less per image (a generous estimate) will result in negligible computational requirements for the classification stages.

In the detection stage, the image was normalized to have a uniform background mean intensity by dividing each pixel by the average of all pixels in the image that were at the same range. Next, the image was divided into three equally sized regions along range to account for the variability of the background and mine signature as a function of range. The image was then convolved with a conventional highlight/shadow matched-filter mask, employing a different mask for each of the three range regions. Next, negative values of the matched-filter image were set to zero. For each range region the resulting matched-filtered image was normalized by removing region mean and dividing by region standard deviation.

A detection was determined by the following process. A target-sized window was scanned over the entire matched-filter image. The number of pixels in the window that exceeded a detection-amplitude threshold were counted. If the count exceeded a detection-count threshold, a detection was declared.

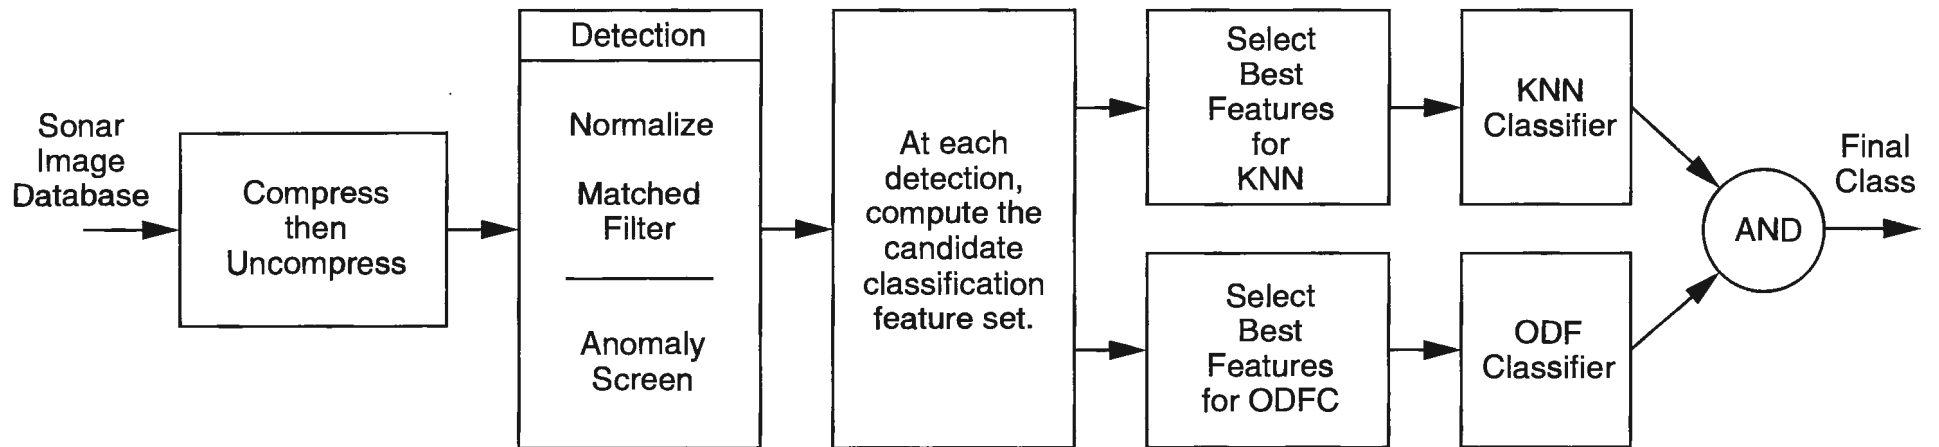


Figure 14. Automatic Algorithm for Detection/Classification of Sea Mines in Sonar Imagery

The optimal detection-amplitude and detection-count threshold pair was selected by a process that computes the optimal receiver operator characteristics (ROC) table. The ROC table contains optimal threshold pairs that give the least FAR for a given Pd.

FEATURE EXTRACTION

In the feature extraction stage, candidate features are computed from a target-sized classification window centered at each detection in both the normalized raw sonar image as well as the matched-filtered image. Forty-five candidate features were computed based on the size, shape strength of the highlight and shadow regions, and on histogram information of pixel intensity within the classification window.

OPTIMAL FEATURE SELECTION

When training classifiers with finite databases, Bellman discovered that the robustness of a classifier collapses when the number of features becomes too large (reference 12). This outcome is known as the curse of dimensionality. Consequently, an important part of the overall classifier design is to select a small and robust set of classification features from the larger candidate set. Evaluating all possible combinations of 45 features is computationally not feasible (there are $2^{45} - 1$ combinations, a number greater than 10^{12}). Therefore, a stepwise optimal selection process was used to rank the features in order of importance. A 45-feature set involves evaluating 1035 feature combinations. Computationally efficient algorithms were developed for this selection process.

Two feature selection algorithms were developed in 1996 by CSS to work specifically with the KNN and ODFC. Bellman's curse of dimensionality implies that for a finite training set, the D/C performance (as judged by the validation/test set) does not continually improve as more features are added (as dimensionality is increased) to the classification process. With a finite data set, it is possible only to determine a subset of features that gives robust performance. Optimal selection means finding the subset of features from a much larger set that gives the best performance for both the finite training and validation (test) sets for a specific classifier (in this study, either KNN or ODFC). Specifically, the best subset is the subset of features that, when fed into the classifier, minimizes a risk metric. The risk metric that is used is defined as "the maximum of the weighted sum of the number of missed mines and number of false alarms."

For a typical candidate feature set of size 45, there are $2^{45} - 1$ combinations of subsets to evaluate in order to determine a global optimum. Because this number is too large to evaluate, a stepwise optimal selection process is used instead of a globally optimum selection process. The stepwise process selects as the first feature the one that gives the best classifier performance of any single feature. The second feature selected is the one that performs best when used with the first. As its name implies, the optimal stepwise selection process continues where at each stage a new feature is selected to add to the subset already determined in previous stages. This stepwise process is obviously not globally optimal, but this concept makes the number of computations feasible and has worked extremely well in this application as it has in other large combinatoric optimization problems.

This research suggests that it is important that the classifier itself be used in the selection process when data sets are finite. Other selection techniques (e.g., principal component analysis or multivariate normal models) are often used because they are computationally tractable. However, they are usually not based on the same mathematical, statistical, or geometrical structure as the classifier; and therefore, in general, they do not select the best features for that

classifier. Because two classifiers are used in this D/C algorithm, there are two selection algorithms. Each one is tuned to the respective classifier because a set of features that is optimal for one classifier is generally far from optimal for the other classifier. This result often occurs when each classifier uses a different mathematical/geometrical structure to partition feature space into class regions. A set of features amenable to one partitioning scheme may not separate well for a different partitioning scheme.

Keeping in mind Bellman's curse of dimensionality, this research suggests that classifiers that train fast, and for which the best subset of features can be optimally selected, will perform significantly better than more sophisticated classifiers for which the best subsets cannot be found because training is too computationally expensive.

CLASSIFICATION

There are two classifiers used in this D/C algorithm: the KNN and the ODFC. Two classifiers that would look at the data through significantly different mathematical, statistical, and geometrical paradigms were picked so that their “ANDING” would be very effective. There could be more than two, but two worked adequately. Also, the process of classifier selection relied strongly on the fact that an optimal feature selection process would be employed that would be tuned to each classifier. With these conditions in mind, the two classifiers were chosen for the following key reasons.

1. They train very fast. This fact makes the algorithm used in the stepwise optimal selection process feasible.
2. They each divide feature space by very different geometrical models—the KNN by hyperspheres, the ODFC by hyperplanes (the term linear combinations is used in reference 8). This process tends to make the classifiers complementary. The ODFC uses a subset of features that can easily separate classes by hyperplanes; the KNN selects a subset that easily separates classes with hyperspheres.
3. The classifiers are somewhat complementary in that they use different noise-tolerant paradigms. The KNN uses sample Bayesian conditional probabilities as determined from the training data; the ODFC uses noise rejection and signal enhancement schemes based on concepts of matched filter design.

KNN is a probabilistic-based neural network that employs radial-basis neurons. During training, K-nearest-neighbor schemes are used to determine proper neuron firing levels. Neurons are conceptualized as hyperspherical attractors in feature space; during training, the number and size of these attractors are determined in an optimal manner that prevents overtraining while covering the entire feature space with enough resolution to determine class boundaries. The training process is noniterative and computationally very fast. The trained network, typically of modest size, executes very quickly.

The ODFC is a classifier with its basis in linear discrimination theory. Two banks of linear filters are determined from the training set. One bank is sensitive to mine characteristics; the other is sensitive to clutter characteristics. The following is a brief description of how the discrimination process works.

Let f be a feature vector. Let $F_m(f, i)$ be the energy output for input f of the i th filter in the mine-filter bank, and let $F_c(f, j)$ be the energy output for input f of the j th filter in the clutter-filter bank. F_m and F_c are linear filters; the outputs of these filters are linear combinations of the input features plus a DC offset and a warping term made up of a linear combination of nonlinear

functions of the features. In the ODFC, linear combinations of features permit boundaries between classes to be described by sets of hyperplanes. By adding linear combinations of nonlinear functions of the features, one changes the hyperplane surface to one that is curved (warped). These nonlinear terms permit the partitioning surfaces in feature space to more efficiently fit irregular class boundaries. Filter coefficients, DC offset, and warping coefficients are determined in such a way that, on the average, the energy output of the mine-filter bank is greater than the energy output of the clutter-filter bank when the input feature vector corresponds to a mine (vice versa for a feature vector that corresponds to clutter); that is, on the average,

$$\begin{aligned} \max_i F_m(f(\text{mine}), i) &> \max_j F_c(f(\text{mine}), j) \\ \max_i F_m(f(\text{clutter}), i) &< \max_j F_c(f(\text{clutter}), j). \end{aligned} \quad (8)$$

The solution involves the solving of a generalized eigenvalue problem.

Classification is determined by ANDING the output of both classifiers. In this process both the KNN and ODFC have two outputs corresponding to the confidence that the input feature vector is associated with a mine or clutter. Let $C_{knn}(\text{mine})$ and $C_{knn}(\text{clutter})$ designate these two confident levels for the KNN, and $C_{odfc}(\text{mine})$ and $C_{odfc}(\text{clutter})$ for the ODFC. If

$$T_{knn} C_{knn}(\text{mine}) > C_{knn}(\text{clutter}), \quad (9)$$

and

$$T_{odfc} C_{odfc}(\text{mine}) > C_{odfc}(\text{clutter}), \quad (10)$$

then a mine classification is declared. The classification gain thresholds, T_{knn} and T_{odfc} are adjusted to select the desired balance between PdPc and FAI.

6. RESULTS

The results of testing the biorthogonal and orthogonal compressed images using the automated classification algorithm are given in table 1. It is interesting to note that the optimal detection thresholds, the optimal classification thresholds, and the optimal feature sets varied as a function of compression method and compression ratio. However, the performance of the D/C algorithm shows that for compression ratios up to 50:1, performance is similar to the no-compression case (which is similar to an expert sonar operator). Even for the 100:1 compression ratio, the performance is quite remarkable. Thus it appears quite feasible to use such compression methods to reduce voluminous amounts of mine countermeasures (MCM) data. Doing so allows real-time communication of sonar data over bandwidth-limited transmission links that are used with remotely operated MCM systems.

Table 1. Results of Testing Using the Automated Classification Algorithm

Compression Ratio (CR)	PdPc %		False Alarms per Image	
	1:1	91		0.28
	Biorthogonal	Orthogonal	Biorthogonal	Orthogonal
25:1	91	91	0.47	0.77
50:1	91	91	0.32	0.72
100:1	91	82	1.50	1.10

7. CONCLUSIONS

The data compression technique has a dual function: (1) compressing the image file size to expedite transmission to the host platform and (2) removing noise and clutter from the image to improve the fidelity of the reconstructed file. For the application at hand (side scan sonar images for detection and classification of underwater mines), the approach selected has been to choose a wavelet-basis and image-decomposition method that best preserves features of underwater mines in the context of side scan sonar images. The results presented show that minimal degradation is realized, even at the highest compression ratios, when the reconstructed images are applied to the automated classification algorithm. Figures 8 through 13 also show that good images are still available for operator evaluation and possible reconciliation of targets and false alarms when the techniques are used in an operator-assist mode. Because of the speed and simplicity of wavelet filters, the biorthogonal and orthogonal data compression techniques offer very capable real-time algorithms for preprocessing side scan sonar images communicated from remotely deployed MCM vehicles to the host platforms through a single communication channel. The results presented in this TM are the first results of this project. Further analysis of the orthogonal and biorthogonal techniques will be pursued to explore methods to improve performance with automated detection and classification algorithms.

REFERENCES

1. G. Wallace, "The JPEG Still Compression Standard," *Communications of the Association of the Computing Machinery*, **34** (4), 1991 pp. 30-44.
2. J. Impagliazzo, W. Greene, and Q. Huynh, "Wavelet Image Compression Algorithm for Side Scan Sonar and Teleradiology," *Proceedings of Wavelet Applications II; SPIE-International Society for Optical Engineering Conference*, 17-21 April 1995.
3. M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, "Image Coding Using Wavelet Transform," *IEEE Transactions on Image Process*, **1** (2), April 1992, pp. 205-220.
4. S. Mallat, "Multiresolution Approximations and Wavelet Orthonormal Bases of $L^2(\mathbb{R})$," *Transactions of the American Mathematical Society*, **315**, 1989, pp. 69-88.
5. I. Daubechies, "Ten Lectures on Wavelets," **61** *Conference Board of the Mathematical Sciences-National Science Foundation Regional Conference Series in Applied Mathematics*, Society for International Applied Mathematics, 1992.
6. R.R. Coifman and M.V. Wickerhouser, "Entropy-Based Algorithms for Best Basis Selection," *IEEE Transactions on Information Theory*, **38** (2), March 1992.
7. R.R. Coifman, Y. Meyer, and M.V. Wickerhouser, "Wavelet Analysis and Signal Processing," in *Wavelets and their Applications*, Ruskai et al., eds., Jones and Bartlett, Boston, MA, 1992.
8. R.W. Hamming, *Coding and Information Theory*, Prentice-Hall, Englewood Cliffs, NJ, 1980.
9. W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B. P. Flannery, *Numerical Recipes in C, The Art of Scientific Computing*, Cambridge University Press, New York, 1992.
10. I.H. Witten, R. Neal, and J.G. Cleary, "Arithmetic Coding for Data Compression," *Communications of the Association of the Computing Machinery* **30** (6), 1987, pp. 520-540.
11. J.M. Shapiro, "Embedded Image Coding Using Zerotrees of Wavelet Coefficients," *IEEE Transactions on Signal Processing*, **41** (12), 1993, pp. 3445-3462.
12. Final Report, Contract DCA100-91-C-0031, Option Year 3; Submitted by: Delta Information Systems, Inc., Horsham, PA, to National Communications System, Arlington, VA, 1994.
13. G.J. Dobeck and J.C. Hyland, "Automated Detection and Classification of Sea Mines in Sonar Imagery," *Proceedings of the SPIE 11th Annual International Symposium on Aerospace/Defense Sensing, Simulation and Control*, April 1997, pp. 21-24.

DISTRIBUTION LIST

Internal

Codes: 10
251
30
30A
38
382
5441 (2)
80
81
811
812
821
8212 (J. Impagliazzo (100))
822
823
831
833

Total: 117