

Detection of Underwater Targets Using a Subspace-Based Method With Learning

Mahmood R. Azimi-Sadjadi, *Senior Member, IEEE*, Nick Klausner¹, *Member, IEEE*, and Justin Kopacz

Abstract—This paper introduces a new subspace-based detection method for multichannel (high frequency and broadband) synthetic aperture sonar (SAS) imagery. An image-dependent dictionary learning method is applied to form the appropriate dictionary matrices for representing target and nontarget image snippets. The hypothesis testing is done by forming a test statistic that relies on the residual error power ratio in representing an unknown image snippet using the target and nontarget dictionary matrices. To avoid the computational bottleneck in most dictionary learning methods, a new recursive method is introduced which does not require any matrix inversion or singular value decomposition (SVD). The proposed detection method was then implemented and benchmarked against a matched subspace detection method for detecting mine-like objects. Results are then presented on two sonar imagery data sets collected in two geographically disparate locations.

Index Terms—Binary hypothesis testing, recursive dictionary learning, sparse coding, synthetic aperture sonar, underwater target detection.

I. INTRODUCTION

DETECTION of underwater objects from sonar imagery is a complicated problem due to numerous factors such as changing operating and environmental conditions; the presence of spatially varying bottom clutter; variations in object highlight and shadow structures with respect to aspect angle, range, and grazing angle; diverse object compositions, and an abundance of nontarget objects that can increase the likelihood of false alarms.

Previous efforts were devoted to the development of both single-channel and multichannel (or multisonar) detection methods [1]–[11]. A basic single-channel approach typically exploits a segmentation algorithm to first isolate the potential contacts and then computes a matching score between a set of common geometrical attributes of the highlight-shadow regions and those of the segmented objects in sonar images to declare a target. A similar approach is also used in [2] though instead of using geometrical relationships, expected highlight-shadow templates

are formed and matched against the extracted objects. More recently, a fast algorithm for the detection of underwater targets in SAS imagery was proposed [3] using a cascaded architecture which operates progressively on smaller portions of the image at each of the three major stages, i.e., shadow detection, ripple detection, and echo detection. This algorithm exploits an integral-image representation for background estimation and generating the shadow map. Although such approaches are successful for high-quality sonar images, their effectiveness remains in doubt when heavily cluttered sonar images are encountered. Additionally, though partially buried targets may not produce prominent highlight-shadow structures, pronounced seafloor clutter (e.g., rocks and coral) could. In such cases, these detectors tend to generate a high incidence of false alarms and missed detections, hence leading to poor performance.

Multichannel detection methods that use multiple sonar frequency bands, e.g., high frequency (HF) and broadband (BB), have addressed some of the shortcomings of the single channel systems. For example, in [10], canonical correlation analysis (CCA) [12] was used to develop a dual-channel (HF and BB) detection method which tests the hypothesis that the presence of a target in a region of interest (ROI) produces a higher level of coherence than when that ROI only contains background clutter. Results in [10] indicated substantially better detection and false alarm rate performance compared to single-channel-based methods. The work in [10] was extended in [11] to multisonar (>2) scenarios using multichannel coherence analysis (MCA) [13], [14]. The expression for the log-likelihood ratio and the J-divergence detectability measure were developed to assess how increasing the number of sonar channels can affect the detection performance. In [15], the problem was posed as a test of independence among the complex beamformed data of multiple sensor channels. More specifically, the test for independence becomes a test of block-diagonal structure in a composite covariance matrix of the channels. It was shown that under the two-dimensional wide-sense stationary assumption (i.e., both along-track and range dimensions) of each channel, the generalized likelihood ratio test (GLRT) leads to a broadband coherence test which can be implemented very efficiently. The detection performance of this method was demonstrated [16] on several multisonar data sets.

The goal of this work is to develop and test a new subspace-based detection method using a fast dictionary learning algorithm. The idea is to perform matching in the sparsely coded domain using image-specific dictionaries to determine if ROIs extracted from pairs of coregistered sonar images contain

Manuscript received August 4, 2015; revised January 31, 2016 and October 2, 2016; accepted December 1, 2016. Date of publication January 11, 2017; date of current version October 11, 2017. This work was supported by the U.S. Office of Naval Research (ONR) under Contracts N00014-12-C-0017 and N00014-15-C-5054.

Associate Editor: T. O. Sæbø

M.R. Azimi-Sadjadi and N. Klausner are with Information System Technologies, Inc., Fort Collins, CO 80525 USA (e-mail: mo@infsyst.biz; nickk@infsyst.com).

J. Kopacz is with the Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO 80523 USA.

Digital Object Identifier 10.1109/JOE.2016.2637138

potential targets. The hypothesis here is that a dictionary formed exclusively from data samples of one class (target or nontarget) will most likely fail to represent data samples from the opposite class. A measure for testing the ability of each dictionary to accurately represent data extracted from each HF–BB ROI pair is then provided. To show the robustness of the detector to the changes in operational and environmental conditions, testing is carried out on two data sets collected in two widely disparate geographical locations. The results are then compared against those of a similar technique that relies on the use of a trained dictionary to detect the presence of targets in the HF and BB images using a matched subspace detector [17].

This paper is organized as follows. Section II describes the proposed dual-channel subspace-based detection method using sparse coding and dictionary learning. Section III gives a brief review of the K-SVD dictionary learning method [18] considered here. A sparse coding algorithm using orthogonal matching pursuit (OMP) [19]–[22] is also reviewed in this section. Furthermore, a new recursive algorithm is introduced for dictionary learning providing the same results as the K-SVD method without any SVD operation. Detection and benchmarking results are presented in Section IV on two sonar imagery data sets. Finally, Section V makes concluding remarks and discusses future lines of research.

II. A SUBSPACE-BASED DETECTION METHOD FOR MULTICHANNEL DATA

In this section, we describe the idea behind the decision criterion used in the proposed subspace-based detection method. This detector operates on the multichannel data (observation) vectors that are obtained by vectorizing and concatenating image snippets or ROIs extracted from every pair of coregistered HF and BB sonar images. More specifically, this composite observation vector $\underline{y}_k = \begin{bmatrix} \underline{y}_k^{(HF)T} & \underline{y}_k^{(BB)T} \end{bmatrix}^T \in \mathbb{R}^N$ contains the magnitude of the beamformed complex-valued data of the k th HF and BB ROIs represented by vectors $\underline{y}_k^{(HF)}$ and $\underline{y}_k^{(BB)}$, respectively, where the superscript T notation denotes the transpose operation. The ROIs overlap in both the range and along-track directions. It is assumed that each composite vector \underline{y}_k can be formed using some linear combination of basis vectors associated with either the null hypothesis, i.e., target absent (or background clutter alone) or the true hypothesis, i.e., target present. That is

$$\begin{aligned} \mathcal{H}_0 : \underline{y}_k &= D_0 \underline{x}_k + \underline{n} \\ \mathcal{H}_1 : \underline{y}_k &= D_1 \underline{x}_k + \underline{n} \end{aligned} \quad (1)$$

where $\underline{x}_k \in \mathbb{R}^M$ is the associated parameter vector with $M \ll N$, $D_0 \in \mathbb{R}^{N \times M}$ and $D_1 \in \mathbb{R}^{N \times M}$ are the subspace or dictionary matrices corresponding to the background clutter and target data, respectively, and \underline{n} is the noise vector that represents the inaccuracy in this representation.

The principal idea here is that the composite observation vector \underline{y}_k lies in a lower dimensional subspace represented by the columns of D_0 or D_1 matrices and \underline{x}_k is a sparse vector with

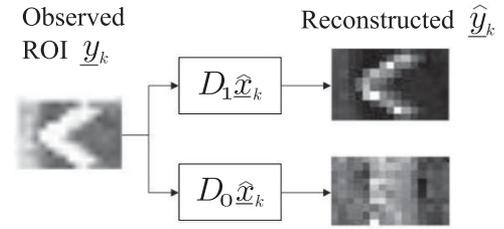


Fig. 1. Exclusive dictionary reconstructions for a target sample.

only a few nonzero elements or atoms that defines this lower dimensional subspace. To use this detector successfully, dictionary matrices must be accurately constructed from the training data sets under each hypothesis. These dictionary matrices can be generated using a subspace construction method such as K-SVD [18], which offers a signal-specific dictionary learning.

Having constructed the subspace matrices D_0 and D_1 based upon some training data, the hypothesis test for an unknown composite observation vector \underline{y}_k can be carried out by first computing the feature vector $\hat{\underline{x}}_k$ using a sparse coding method [23] under each hypothesis and then using the decision rule

$$m^* = \arg \min_{m \in \{0,1\}} \|\underline{y}_k - D_m \hat{\underline{x}}_{k,m}\|^2 \quad (2)$$

where $\hat{\underline{x}}_{k,m}$ is the estimate of \underline{x}_k generated using the sparse coding when dictionary D_m was used. The idea here is that while each dictionary matrix is suitable in representing its specific data, it will poorly represent data from the opposite hypothesis. That is, if one attempts to reconstruct an observation containing a target using the background clutter dictionary matrix one will produce a large residual, while the reconstruction of a target will be accurate for the target dictionary matrix and *vice versa*. This idea is depicted in Fig. 1, where the input ROI containing a target is sparsely represented as using dictionaries D_1 and D_0 . From the reconstructed images $\hat{\underline{y}}_{k,1} = D_1 \hat{\underline{x}}_{k,1}$ and $\hat{\underline{y}}_{k,0} = D_0 \hat{\underline{x}}_{k,0}$, it is clear that the input image is most similar to the image that was reconstructed using atoms taken from the target dictionary D_1 . This similarity can be represented mathematically by simply looking at the norm of the reconstruction error, which is lower for the dictionary containing atoms representing ROIs of the same class as the input, i.e., D_1 in this case.

Alternatively, to perform the decision making on the composite observation vector \underline{y}_k , we form the residual ratio $\Lambda(\underline{y}_k)$, and compare it to a prespecified threshold γ

$$\Lambda(\underline{y}_k) = \frac{\|\underline{y}_k - D_0 \hat{\underline{x}}_{k,0}\|_2^2}{\|\underline{y}_k - D_1 \hat{\underline{x}}_{k,1}\|_2^2} \underset{\text{Background}}{\overset{\text{Target}}{\geq}} \gamma. \quad (3)$$

The threshold γ is chosen to achieve a desired false alarm rate. This choice is heavily dependent upon the environment and background clutter density and how well the dictionaries are able to represent the data in each ROI. The entire detection process is shown in Fig. 2 where a pair of coregistered ROIs forming the composite test vector \underline{y}_k is applied to the system. This vector is then sparsely coded (with OMP) using two dictionary matrices

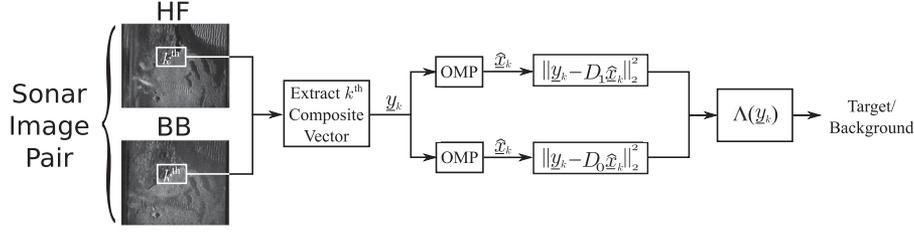


Fig. 2. Subspace-based detection process.

D_1 and D_0 for target and nontarget, respectively. The residual ratio $\Lambda(\underline{y}_k)$ is then compared against the chosen threshold for decision making.

The decision making in this detector requires constructing an optimum image-dependent dictionary matrix D_m for each hypothesis $m = 0, 1$. In the next section, we give a brief review of the theory behind the K-SVD method and our new and more efficient dictionary learning and sparse coding methods.

III. DICTIONARY LEARNING AND SPARSE CODING

The name K-SVD is synonymous with the K-means algorithm that updates codebook vectors for clustering and vector quantization. Similarly, K-SVD obtains an updated dictionary basis by performing an SVD operation at every step of the learning. Under certain special cases, K-SVD produces the same results as K-means [18]. The K-SVD method has been applied to many sparse-representation-based pattern recognition problems. For example, in [24], Zhang and Li developed a discriminative K-SVD (D-KSVD) by incorporating the classification error into the objective function.

The main purpose of K-SVD is to create a signal-specific dictionary that can be used to reduce the dimension of an observation vector by representing it as a sparse linear combination of relatively few atoms. K-SVD achieves this by solving a constrained optimization problem to minimize the reconstruction error in a set of training vectors. Let $Y \in \mathbb{R}^{N \times K}$ be a matrix containing the training data vectors \underline{y}_k , $k \in [1, K]$ as its columns, $D \in \mathbb{R}^{N \times M}$ be a signal-dependent dictionary matrix to be found, and $X \in \mathbb{R}^{M \times K}$ be the sparse representation of Y in the subspaces spanned by the columns of D . Note that here we drop the hypothesis index m for simplicity in notation. However, for each hypothesis, the matrix D_m must be found based upon its own training data matrix Y_m . Also, the number of nonzero elements of each \underline{x}_k must be substantially less than N to enforce sparsity. This constrained optimization problem is then given [18] by

$$\min_{D, X} \|Y - DX\|_F^2 \quad \text{s.t.} \quad \|\underline{x}_k\|_0 \leq \tau \quad \forall k \quad (4)$$

where $\|A\|_F^2 = \text{tr}(AA^H)$ is the Frobenius norm of matrix A , $\|\cdot\|_0$ denotes the l_0 -norm which is the number of nonzero elements of a vector, and τ is the sparsity factor.

During the training, the K-SVD algorithm iterates between two phases: 1) a sparse representation phase where, for each training data vector \underline{y}_k , the corresponding reduced dimensional

vector \underline{x}_k is computed based on a given D at the previous iteration using some pursuit method, e.g., OMP; and 2) a dictionary update phase where D is updated in the next iteration based on minimizing the reconstruction error using SVD. In what follows, summaries of these two steps are provided.

A. Sparse Coding Phase

In the sparse coding phase, the goal is to minimize the cost function in (4) or alternatively $\sum_{k=1}^K \|\underline{y}_k - D\underline{x}_k\|_2^2$ by finding sparse solution vectors \underline{x}_k 's, given the data matrix Y and the current dictionary matrix D . Although any pursuit algorithm [19]–[22] can be employed for this sparse coding, here we have used the OMP method.

The OMP algorithm is a greedy method which iteratively selects the best dictionary atoms for each \underline{y}_k to reduce the reconstruction error (or residual) at each iteration. The algorithm starts by initializing the reconstruction error as the observation vector $\underline{r}_0 = \underline{y}_k$ and the set of selected basis as $S_0 = \emptyset$, i.e., the empty set. At every iteration t , the residual vector \underline{r}_{t-1} is projected onto all the unselected dictionary basis vectors and the one which yields the largest inner product (i.e., the best chance to reduce the error) is selected, and the index set is updated. That is, $S_t = S_{t-1} \cup \{l_t\}$, where

$$l_t = \underset{j}{\text{argmax}} \quad |\underline{r}_{t-1}^T \underline{d}_j| \quad \forall j \notin S_{t-1} \quad (5)$$

and \underline{d}_j is the j th column of matrix D or dictionary basis vector. An augmented dictionary of selected basis is then created to reflect the new index in S_t by including the new basis vector, i.e., $D^{(t)} = [D^{(t-1)} \quad \underline{d}_{l_t}]$. After updating $D^{(t)}$, the new sparse signal $\hat{\underline{x}}_k^{(t)}$ can be found using the least squares (LS) method as

$$\hat{\underline{x}}_k^{(t)} = [D^{(t)T} D^{(t)}]^{-1} D^{(t)T} \underline{y}_k = Q_{D^{(t)}} \underline{y}_k \quad (6)$$

where $Q_{D^{(t)}}$ represents an LS filtering operator or the Moore–Penrose inverse matrix of $D^{(t)}$. Using this sparse vector $\hat{\underline{x}}_k^{(t)}$, the new residual must be computed as

$$\underline{r}_t = \underline{y}_k - D^{(t)} \hat{\underline{x}}_k^{(t)}. \quad (7)$$

If we define the projection matrix onto the subspace spanned by the columns of $D^{(t)}$ by $P_{D^{(t)}} = D^{(t)} (D^{(t)T} D^{(t)})^{-1} D^{(t)T}$, then from a geometrical perspective the residual \underline{r}_t can be viewed simply as the projection of \underline{y}_k onto the orthogonal complement of subspace $\langle D^{(t)} \rangle$, i.e., $\underline{r}_t = P_{D^{(t)}}^\perp \underline{y}_k$ where $P_{D^{(t)}}^\perp = I - P_{D^{(t)}}$. Note that we have $Q_{D^{(t)}} P_{D^{(t)}} = Q_{D^{(t)}}$. A

stopping criterion can be chosen either by reaching a predefined sparsity limit or when the magnitude of the residual vector falls below a threshold.

As can be seen from (6), the computational effort of the original OMP algorithm grows very fast as the number of atoms increases. This is mainly due to the matrix inversion operation in (6) as the dimension of matrix $[D^{(t)T} D^{(t)}]$ increases with the number of atoms selected. In the Appendix, we provide a summary of the fast OMP algorithm in [25] which avoids using any matrix inversion operation.

B. Dictionary Update Phase

For the dictionary update phase [18], the cost function in (4) is rewritten as

$$\begin{aligned} \|Y - DX\|_F^2 &= \left\| Y - \sum_{j=1}^M \underline{d}_j \underline{x}_T^j \right\|_F^2 \\ &= \left\| \left(Y - \sum_{j \neq k} \underline{d}_j \underline{x}_T^j \right) - \underline{d}_k \underline{x}_T^k \right\|_F^2 \\ &= \|E_k - \underline{d}_k \underline{x}_T^k\|_F^2 \end{aligned} \quad (8)$$

where \underline{x}_T^k is the k th row vector of matrix X formed of the all the coefficients corresponding to k th atom, and E_k is the error matrix which represents the reconstruction error when excluding the contribution of the k th atom. In [18], the SVD was used to find \underline{d}_k that minimizes the error in (8). Below we provide a brief summary of the steps using the same notations as used in [18].

To impose the sparsity condition, we let $\omega_k = \{i | x_T^k(i) \neq 0\}$ be a subset of indices corresponding to the training samples in Y that use the dictionary atom \underline{d}_k , i.e., those with $x_T^k(i) \neq 0$ where $x_T^k(i)$ is the i th element of \underline{x}_T^k . Clearly, $Q = |\omega_k| \leq K$ where $|\omega_k|$ is the cardinality of set ω_k . Now, if we define matrix Ω_k of size $K \times Q$ whose $(\omega_k(i), i)$ entries are ones and zeroes elsewhere, then $\underline{x}_R^k = \underline{x}_T^k \Omega_k$ is the restricted and reduced dimensional row vector version of \underline{x}_T^k with all zero entries removed. Also, $E_k^R = E_k \Omega_k$ is the restricted error matrix that contains only Q columns of E_k corresponding to those samples that use \underline{d}_k . The cost function in (8) can then be rewritten as

$$\|E_k \Omega_k - \underline{d}_k \underline{x}_T^k \Omega_k\|_F^2 = \|E_k^R - \underline{d}_k \underline{x}_R^k\|_F^2. \quad (9)$$

Now, the SVD can be used to decompose $E_k^R = U \Delta V^T$ to give the solution for the updated dictionary atom $\hat{\underline{d}}_k$. This best rank-one solution [26] corresponds to the left singular vector corresponding to the largest eigenvalue, i.e., $\hat{\underline{d}}_k = \underline{u}_1$ where \underline{u}_1 is the first column of U .

Because $\hat{\underline{d}}_k$ is changed, \underline{x}_R^k must also be updated to reflect these changes, otherwise E_k would be computed incorrectly for the next atom. The restricted vector \underline{x}_R^k is updated using $\hat{\underline{x}}_R^k = \underline{v}_1 \Delta(1, 1)$ where \underline{v}_1 is the first row of V^T , and $\Delta(1, 1)$ is the first diagonal element of matrix Δ . After updating the k th dictionary atom, the same procedure is repeated for the next atom until all atoms have been individually updated. Using this

newly updated dictionary matrix, the sparse matrix X can be recomputed using the OMP method. The dictionary update and sparse representation phases will be repeated until the stopping conditions are met (either maximum number of iterations or a reconstruction error below some chosen threshold).

As can be seen, the solution of \underline{d}_k is obtained by applying the SVD to E_k^R and choosing the first left singular vector. However, this process is computationally inefficient especially for large dimensional matrices as SVD returns all the left and right singular vectors and values, though we are only interested in the most dominant ones. In the following section, we introduce a new recursive algorithm to arrive at the solution for \underline{d}_k without using any SVD operation.

C. Recursive Updating for Dictionary Learning

To arrive at an efficient algorithm that avoids any SVD operation, we pose the problem in (9) as a constrained minimization

$$\min_{\underline{d}_k} \|E_k^R - \underline{d}_k \underline{x}_R^k\|_F^2 \quad \text{s.t.} \quad \|\underline{d}_k\| = 1 \quad (10)$$

where the constraint is added to impose the unit norm requirement for \underline{d}_k . The Lagrangian associated with (10)

$$L(\underline{d}_k, \lambda) = \|E_k^R - \underline{d}_k \underline{x}_R^k\|_F^2 + \lambda(\underline{d}_k^T \underline{d}_k - 1) \quad (11)$$

has the following first-order partial derivatives:

$$\frac{\partial L(\underline{d}_k, \lambda)}{\partial \underline{d}_k} = 2 \left(\underline{x}_R^k \underline{x}_R^{kT} + \lambda \right) - 2E_k^R \underline{x}_R^{kT} \quad (12)$$

$$\frac{\partial L(\underline{d}_k, \lambda)}{\partial \lambda} = \underline{d}_k^T \underline{d}_k - 1. \quad (13)$$

Setting (12) and (13) to zero and solving this system of equations for \underline{d}_k , we obtain

$$\underline{d}_k = \frac{E_k^R \underline{x}_R^{kT}}{\sqrt{\underline{x}_R^k E_k^{RT} E_k^R \underline{x}_R^{kT}}}. \quad (14)$$

This solution requires having previously computed \underline{x}_R^k through OMP. Since the restricted vector \underline{x}_R^k changes when \underline{d}_k is updated, we also need to update \underline{x}_R^k . This can easily be done by taking the partial derivative of the cost function in (9) with respect to \underline{x}_R^k and solving for \underline{x}_R^k

$$\underline{x}_R^k = \frac{\underline{d}_k^T E_k^R}{\underline{d}_k^T \underline{d}_k}. \quad (15)$$

The coupled equations in (14) and (15) must be iterated until convergence. To start the process, we use \underline{x}_R^k that has been computed by OMP, as the entire sparse matrix X is available after the sparse pursuit phase, to initialize \underline{d}_k with (14). After initializing \underline{d}_k , \underline{x}_R^k can be updated for the first time using (15). It can easily be shown that the convergence is assured and the reconstruction error will monotonically decrease as the dictionary update takes place.

To give a simple illustration, the algorithm was used to recursively extract the singular vectors of the training data set Y_1 consisting of 16 targets used to construct the dictionary D_1 in the detection studies of Section IV. That is, the expressions

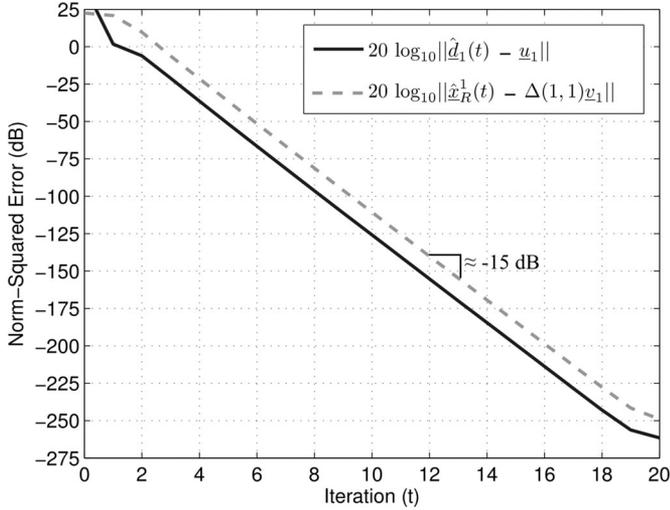


Fig. 3. Convergence of vectors $\hat{d}_1(t)$ and $\hat{x}_R^1(t)$ to their singular vectors \underline{u}_1 and $\Delta(1, 1)\underline{v}_1$, respectively.

given in (14) and (15) were computed in an alternating fashion but with the matrix Y_1 replacing the error matrix E_k^R in both equations. At each iteration t , the pair of vectors $\hat{d}_1(t)$ in (14) and $\hat{x}_R^1(t)$ in (15) were compared to their corresponding singular vectors \underline{u}_1 and $\Delta(1, 1)\underline{v}_1$ found by computing a full SVD of Y_1 as discussed in Section III-B. Fig. 3 plots the norm-squared error in both vectors as a function of the iteration index t . As can be seen from this plot, both vectors rapidly converge to their corresponding singular vectors, achieving an error below -100 dB (or a norm-squared error of 1×10^{-10}) after only ten iterations. As annotated in the plot, the algorithm tends to produce a linear decrease in error at a rate of approximately 15 dB per iteration.

The order of computations for performing both recursive equations in (14) and (15) is $\mathcal{O}(NQ)$ where $Q = |\omega_k| \ll N$ whereas for the SVD operation, the order of computations is approximately $\mathcal{O}(N^3)$. This iterative algorithm, coupled with the fast OMP in the Appendix, provides an efficient scheme for dictionary learning and sparse coding free of any matrix inversion and SVD operation.

IV. DETECTION RESULTS ON SONAR IMAGERY

In this section, we provide the results of the proposed subspace-based detector on two small synthetic aperture minehunter (SSAM) data sets [27]. A more detailed description of each data set and their properties is given in Section IV-A. The SSAM system collects pairs of coregistered HF and BB sonar images for the purposes of underwater target detection and classification. The HF sonar images provide better spatial resolution to capture target details and characteristics, i.e., the highlight and shadow structures typically have good definitions with each having a definite shape and adequate number of pixels. Nonetheless, these highlight-shadow structures are also prominent for structured clutter such as coral reefs, rocks, and sand ripples making automatic detection and classification difficult. The BB sonar images, on the other hand, offer better clutter

suppression, though a prominent shadow structure may not be present for some targets. Therefore, the use of dual frequency sonar systems allows one to exploit a high-resolution HF sonar with good target definition while taking advantage of the clutter suppression ability of a low-resolution BB sonar coregistered over the same region to provide better detection performance comparing to those of the single-sonar cases.

The process begins by partitioning each pair of coregistered HF and BB sonar images into ROIs of size 80×140 pixels by sweeping a window through each image. The ROIs are shifted by only 10 pixels (corresponding to approximately 0.25 m) in both the range and along-track directions to ensure that the target will be approximately centered in at least one of the ROIs. The magnitude of the beamformed complex-valued data of every pair of HF and BB ROIs is vectorized and concatenated to form the composite vector \underline{y}_k used for training or testing purposes.

The training phase involves finding dictionary matrices D_m $m = 0, 1$ for target and nontarget cases using a small subset of the samples taken from the first data set. Samples from data set 2 were not included in order to test the effectiveness of the detector on a totally new data set collected in a geographically disparate location. Fig. 4 shows the selected training samples (only HF ROI are shown) for 16 targets and 30 background clutter. An attempt was made to select a representative set of target and background samples from data set 1 for the purpose of training. The dictionary learning algorithm in Section III-C was applied to these training sets to provide minimum reconstruction error [in the minimum mean square error (MMSE) sense] while reducing the number of components or atoms needed for such reconstruction via the sparse coding by imposing a sparsity constraint of $\tau = 13$. Two dictionary matrices D_1 and D_0 were created using their respective training data matrices Y_1 of size 22400×16 and Y_0 of size 22400×30 that contain only target and background clutter (nontarget) samples, respectively. Fig. 4 also shows five (out of 13) corresponding atoms for both target and nontarget dictionaries. As expected, the target dictionary D_1 contains atoms that capture different distinct features of the selected target samples whereas the atoms in D_0 appear to be a mixture of random noise and odd patterns of rocks and coral. This is due to the selected nontarget training images. The displayed atoms are those that are needed to minimize the reconstruction error for the respective training sets. Once these dictionary matrices are obtained, the subspace detector can then be applied to the testing data sets.

To study the algorithm's ability to sparsely represent the targets in these two data sets, snippets of each target from both data sets were extracted. Using the dictionary matrix D_1 found using the aforementioned procedure, the ratio

$$\eta = \frac{\|\underline{y}_k - D_1 \hat{x}_k\|^2}{\|\underline{y}_k\|^2} \quad (16)$$

was computed as a function of the number of nonzero values in \hat{x}_k estimated using the OMP algorithm in the Appendix. This was done for each \underline{y}_k corresponding to a target for both data sets. The ratio η in (16) measures the ability of dictionary D_1 to sparsely represent the targets in these two data sets as a function

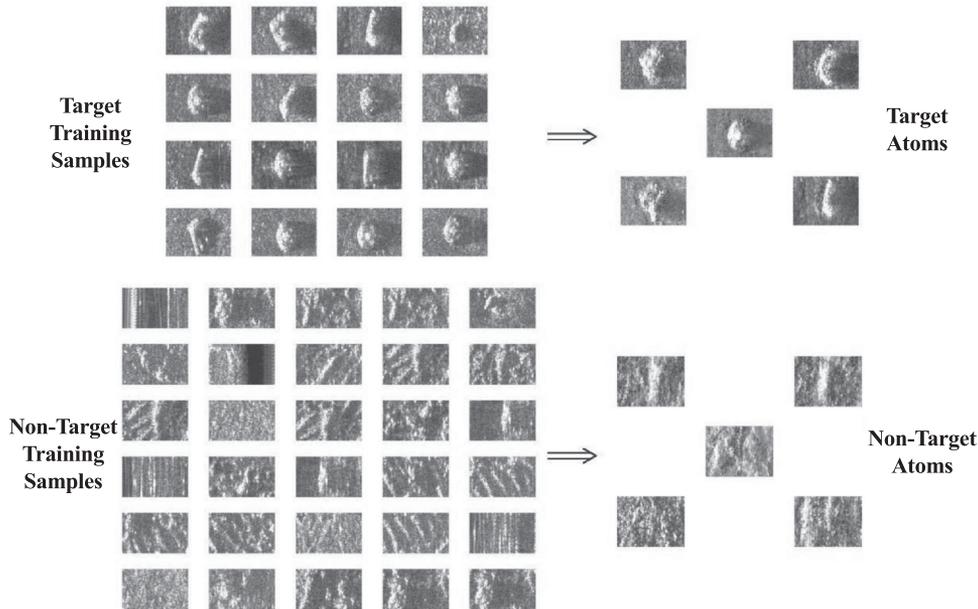


Fig. 4. Training samples and the corresponding dictionary atoms.

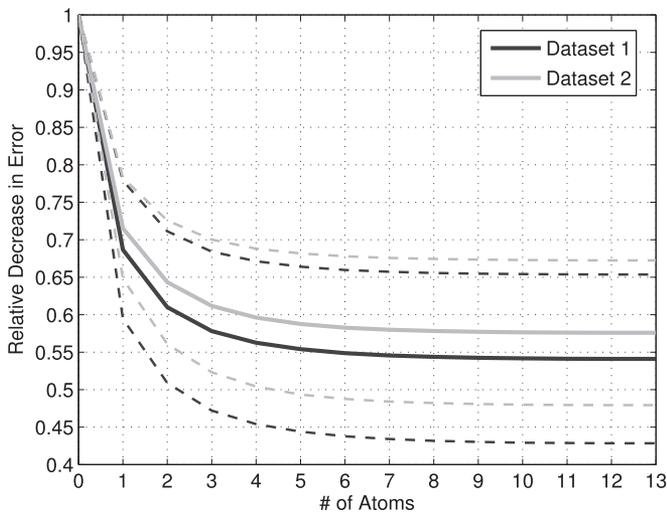


Fig. 5. Relative decrease in error versus the number of atoms for the targets from both data sets.

of the number of atoms used in that dictionary relative to the total energy in the observation vector \underline{y}_k . The dark and light solid lines in Fig. 5 plot the average value of η for targets in data set 1 and data set 2, respectively. The dashed lines above and below each corresponding line plot one standard deviation from the mean. From the results of this figure, one can see that the relative reconstruction error tends to be lower for the targets of data set 1 compared to those in data set 2. This is to be expected given that the dictionary D_1 was trained exclusively from a subset of targets in data set 1. More importantly, however, one can see that the average relative reconstruction error is at or below 0.6 after using only four to five atoms. Thus, one can see that the proposed method is capable of producing a dictionary that is capable of sparsely representing the targets in these two data sets.

In what follows, we describe the characteristics of the two datasets used to test the detector's performance in two significantly different environmental and operating conditions.

A. SSAM I Data Set Description

The first data set used in this study consists of 122 pairs of HF and BB SAS images containing in total 77 targets of varying sizes, shapes, and ranges from the platform. The clutter density ranges from medium to difficult as images can often contain things such as small sand ripples and large regions consisting of a rocky or coral background. As mentioned before, the dictionary matrices were generated using a training set consisting of a subset of 16 targets as well as a selected set of 30 background ROIs taken from this data set.

The second data set used in this study, on the other hand, is composed of 458 pairs of HF and BB SAS images containing only 68 targets of varying sizes, shapes, and ranges from the sonar platform. In this data set, the amount of targets is vastly outweighed by the amount of clutter giving the added benefit of simulating a wide scale search. This data set was used to test the ability of the subspace-based detector to successfully detect contacts in a significantly different environmental condition. This is because these two data sets were collected off the shores of two different continents. Comparing to data set 1, data set 2 was collected over three different environments adding an extra level of difficulty. Two of the environments were relatively flat with bottom features somewhat similar to those observed in data set 1, whereas the third environment is very complex as it contains a mixture of sea grass and sand formations. The particular species of sea grass is unique to this region, and it is an element foreign to data set 1. Additionally, data set 2 contained some target types that were not present in data set 1. Consequently, it is expected that the dictionary matrices

generated based upon the limited training data taken from data set 1 will have difficulty representing these images.

B. Benchmark Methods

In addition to evaluating the proposed subspace-based detector's ability to detect potential targets, its performance was also compared to a matched subspace detector [17] which is somewhat similar to the proposed method given in (3) in that it relies on the use of a trained dictionary to detect the presence of targets in both sonar images. For this method, it is assumed that the composite observation y_k in (1) can be expressed as

$$y_k = H\theta_k + \mathbf{1}\phi_k + w_k \quad (17)$$

where $\theta_k \in \mathbb{R}^M$ and $\phi_k \in \mathbb{R}$ are deterministic but unknown quantities that parameterize the model, $H \in \mathbb{R}^{N \times M}$ is a fixed dictionary designed to represent targets in both the HF and BB sonar data, $\mathbf{1} \in \mathbb{R}^N$ is an all-one vector included to model an unknown bias in the data, and $w_k \in \mathbb{R}^N$ is a vector of white noise of unknown variance σ^2 . Using the same set of 16 target ROIs used to train the subspace-based detector, the target dictionary H in (17) was formed by taking an SVD [28] of the training data and retaining the same number basis vectors corresponding to the top 13 singular values. In the linear model given in (17), the parameter vector θ_k determines what contribution this target dictionary makes to the observed data and the parameter ϕ_k is included to model an unknown average pixel intensity in the magnitude of the image data due to the background. Using this model, the GLRT [17] is then used to test the null hypothesis that a target is not present ($\theta_k = \mathbf{0}$) versus the alternative that it is ($\theta_k \neq \mathbf{0}$). Comparing this detection strategy to that of the K-SVD-based method in (3), one can see that, while both methods rely on the use of a target dictionary trained using the same exemplar images, the expression in (17) models the background within the snippet as a simple unknown bias whereas the proposed subspace-based detector in (3) relies on a nontarget dictionary D_0 to model the background.

C. Detection Results—Data Set 1

The bold solid and dashed lines in Fig. 6 give the receiver operating characteristic (ROC) curves for the proposed (labeled “subspace-based” in the figure) and matched subspace (labeled “benchmark”) detectors, respectively, when applied to data set 1. Note that in these plots the horizontal axis is the number of false alarm/image. The lighter solid and dashed lines surrounding these two curves represent a 95% confidence interval [29] for each ROC curve. These confidence intervals are based on the assumption that the measured probability of detection and probability of false alarm are estimates of an asymptotically normal random variable whose variance depends on the number of target and nontarget samples used to estimate the ROC curve. The confidence interval may then be constructed by finding the bound that captures 95% of the area underneath that normal density. From this figure, one can see that both detection methods perform equally well on this data set which is to be expected given that both were trained from data taken from this data set. The probability of detection at the knee point

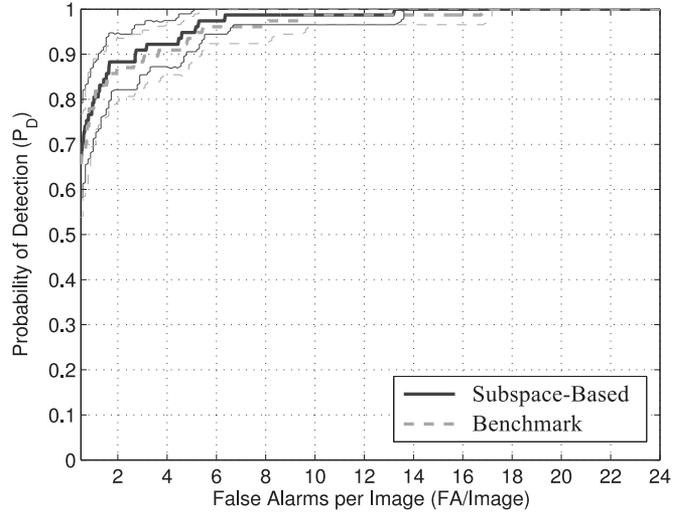


Fig. 6. ROC curves comparing the subspace-based detector and benchmark method for data set 1.

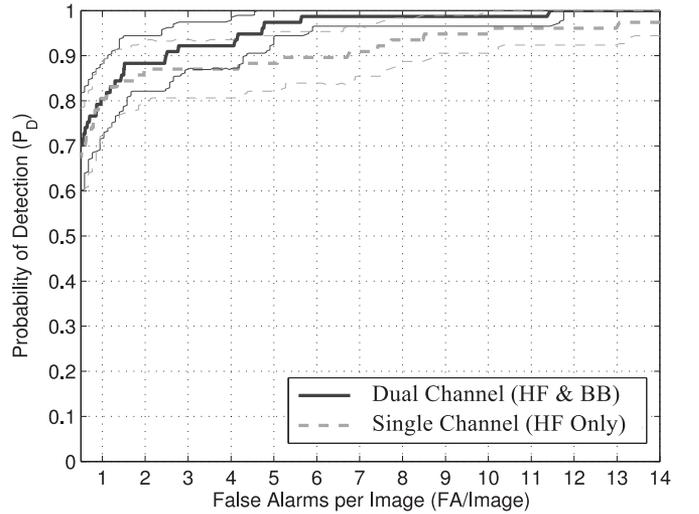


Fig. 7. ROC curves comparing the single and dual channel subspace-based detectors for data set 1.

(where $P_D + P_{FA} = 1$) were $P_D = 0.90$ and $P_D = 0.91$ for the matched subspace detector and the proposed subspace-based detector, respectively.

To observe if and what improvement adding the BB image has on the performance of the proposed detector, Fig. 7 compares the same dual-channel results given in Fig. 6 with its single-channel version, i.e., where the observation vector y_k consists of the HF image alone. The results of this figure seem to suggest that, at low false alarm rates, both the single-channel and dual-channel variants of the subspace-based detector perform similarly. However, at higher false alarm rates the dual-channel subspace-based detector seems to perform better than its single-channel counterpart. This may be due in part to the clutter suppression abilities of the BB sonar leading to a lower false alarm rate at a given probability of detection. At the knee point of the ROC curve, the single-channel detector achieves a probability of detection of $P_D = 0.87$ versus $P_D = 0.91$ for the dual-channel one.

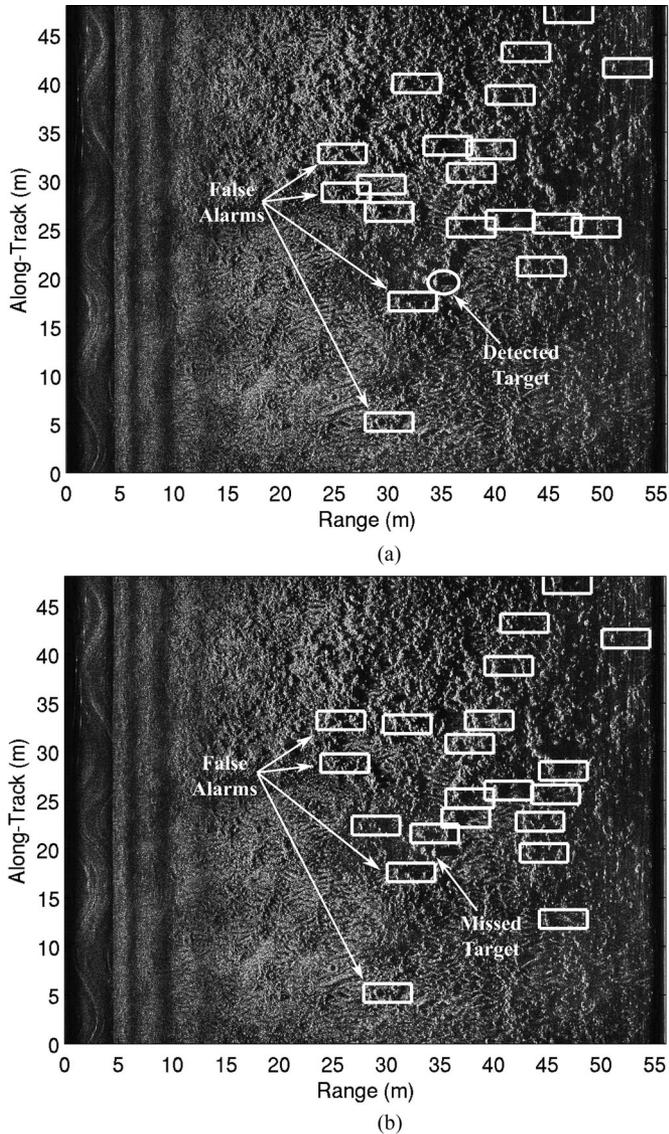


Fig. 8. A difficult sonar image of data set 1 with detected contacts for both detection methods. (a) Subspace-based detector. (b) Matched subspace detector.

Fig. 8(a) and (b) shows the detection results of the subspace-based and matched subspace detectors, respectively, on a typical “difficult” (high density of clutter) HF sonar image of data set 1. As indicated within each image, each rectangular box within both of these figures indicates the location of a false alarm while the locations of detected targets are indicated by a circle. For both of these images, the threshold needed to achieve a false alarm rate of $P_{FA} = 0.01$ was selected for each method. An unsupervised density-based clustering algorithm known as DBSCAN [30] was applied to the set of detection coordinates obtained using both methods to group them into a specific cluster. The ROI with the largest detection statistic within a given cluster can then be used to represent that cluster. This clustering procedure not only reduces false alarms, but also ensures that multiple detected ROIs within close proximity of a target are not declared as separate detections. These results indicate that although both methods exhibit similar false alarm rates for this

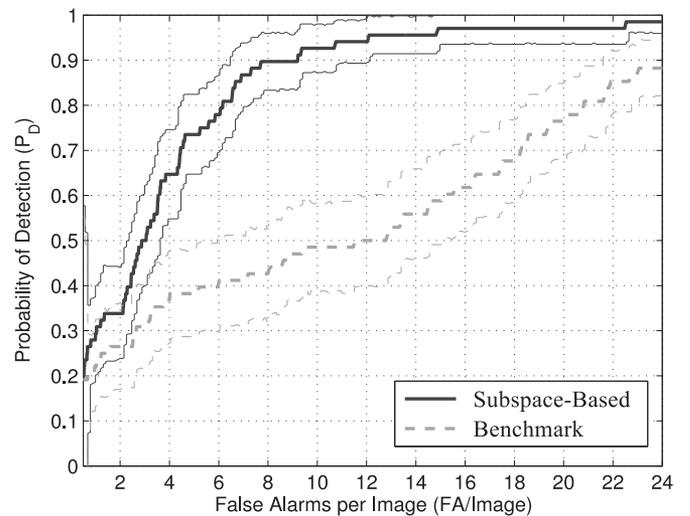


Fig. 9. ROC curves comparing the subspace-based detector and benchmark method for data set 2.

image, the matched subspace detector misses the only target in this case.

D. Detection Results—Data Set 2

The proposed subspace-based detector designed using exemplars from data set 1 was then tested on data set 2 which was collected in a totally different environment in order to test the robustness of the system. Again, the results of the detector are benchmarked against those of the matched subspace detector in [17]. Fig. 9 shows the ROC curves for both methods as well as their 95% confidence intervals. As can be seen, both detectors exhibit a degradation in performance compared to the results of data set 1 in Fig. 6. This is to be expected given that data set 2 presents a completely different set of environments with a completely new set of targets. However, for this data set, one can see that the subspace-based detector brings a significant improvement in detection performance compared to the matched subspace detector, achieving detection probabilities of $P_D = 0.81$ and $P_D = 0.59$ at their respective knee points. Using the same thresholds that were used to generate Fig. 8, Fig. 10(a) and (b) shows the detection results of these detectors on a typical “difficult” HF sonar image of data set 2. Once again, rectangular boxes denote the locations of false alarms within each of these images while detected targets are indicated by a circle. One can see from these two figures that both methods are able to detect the target in the image. However, the subspace-based detector exhibits fewer false alarms compared to its matched subspace counterpart.

Finally, Fig. 11 compares the dual-channel subspace-based detector with its single-channel counterpart for this data set. In contrast to the results of Fig. 7, however, the single-channel version performs just as well if not slightly better than its dual-channel counterpart with a knee-point probability of $P_D = 0.84$. This lack of clear improvement when going from the HF image alone to using both the HF and BB images for this data set may

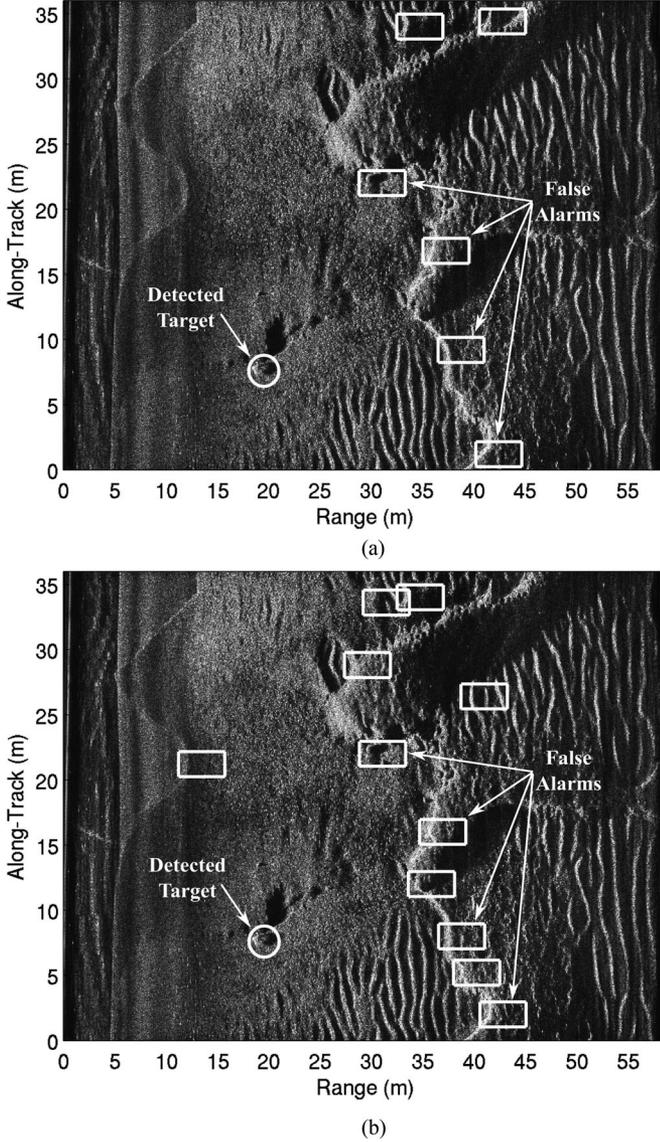


Fig. 10. A difficult sonar image of data set 2 with detected contacts for both detection methods. (a) Subspace-based detector. (b) Matched subspace detector.

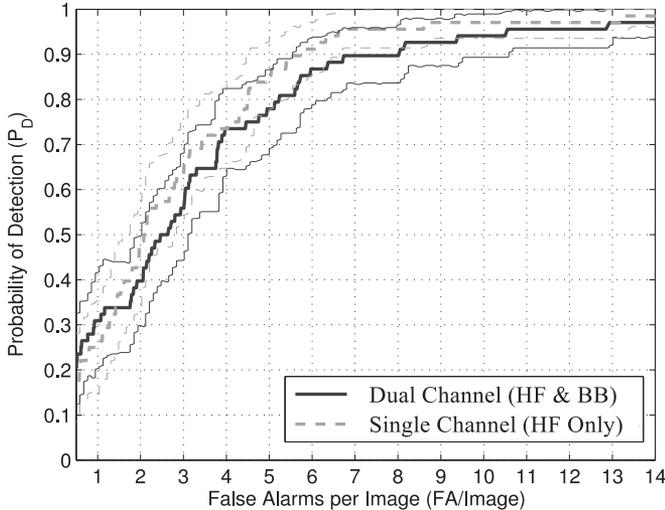


Fig. 11. ROC curves comparing the single- and dual-channel subspace detectors for data set 2.

be due in part to the relatively poor quality of the BB images compared to those of data set 1. Thus, the usefulness of the BB image into this subspace-based detection approach may be something that depends on the environment in which data are collected as well as the sonar system.

V. CONCLUSION

The problem of underwater target detection from dual-channel sonar images was considered in this paper. A new subspace-based detection method using a dictionary learning was introduced which relies on generating a residual ratio based upon the reconstruction error of a particular unknown sample projected onto two different subspaces: one for target and one for nontarget/clutter. These two dictionaries are constructed by applying a new recursive dictionary learning algorithm in conjunction with a fast OMP method to a set of target and nontarget training patterns which yield dictionaries that are optimal at sparsely representing their respective data type. The detector is then applied to two sonar imagery data sets consisting of pairs of HF and BB images coregistered over the seafloor and its performance was benchmarked against a matched subspace detector that similarly relies on a trained target template to identify targets in the image. The results of this study indicated that the proposed subspace-based detection method is much better suited for detecting potential target ROIs in sonar images. However, it was observed that both detection methods exhibited poorer performance when applied to the novel data set 2 which was not included in the training. In addition to being a more difficult and widely different data set, this diminishing performance is most likely attributed to the fact that neither method was able to accurately represent both HF and BB channels of the sonar image owing to substantial differences between the BB sonar images in the first data set used for training and those of the second data set. Future research could include developing an *in situ* dictionary updating method to account for substantial variations in the environmental conditions.

APPENDIX

FAST OMP ALGORITHM

To develop our fast OMP, we use the orthogonal projection updating equations [31] for matrices $P_{D^{(t)}}$ and $Q_{D^{(t)}}$. The projection matrix $P_{D^{(t-1)}} = D^{(t-1)}(D^{(t-1)T}D^{(t-1)})^{-1}D^{(t-1)T}$ at iteration $t-1$ can be updated by appending the new atom \underline{d}_t into the column span of $D^{(t-1)}$. It can be shown [25] that

$$P_{D^{(t)}} = P_{D^{(t-1)}} + P_{D^{(t-1)}}^\perp \underline{d}_t (\underline{d}_t^T P_{D^{(t-1)}}^\perp \underline{d}_t)^{-1} \underline{d}_t^T P_{D^{(t-1)}}^\perp. \quad (\text{A.1})$$

The term $\tilde{\underline{d}}_t = P_{D^{(t-1)}}^\perp \underline{d}_t$ is the projection of \underline{d}_t onto the orthogonal component of subspace $\langle D^{(t-1)} \rangle$ which represents the new part (or innovation) of \underline{d}_t . Thus, (A.1) can be rewritten as

$$P_{D^{(t)}} = P_{D^{(t-1)}} + P_{\tilde{\underline{d}}_t} \quad (\text{A.2})$$

where $P_{\tilde{\underline{d}}_t} = \tilde{\underline{d}}_t / (\tilde{\underline{d}}_t^T \tilde{\underline{d}}_t)$ is the projection matrix for $\langle \tilde{\underline{d}}_t \rangle$. Using (A.2), we also have

$$P_{D^{(t)}}^\perp = P_{D^{(t-1)}}^\perp - P_{\tilde{\underline{d}}_t}^\perp. \quad (\text{A.3})$$

It can also be shown [31], [25] that for $Q_{D^{(t)}}$ we get the following update equation:

$$Q_{D^{(t)}} = \begin{bmatrix} Q_{D^{(t-1)}} & \\ & 0 \end{bmatrix} - \begin{bmatrix} Q_{D^{(t-1)}} \underline{d}_{l_t} \\ -1 \end{bmatrix} (\underline{d}_{l_t}^T P_{D^{(t-1)}}^\perp \underline{d}_{l_t})^{-1} \times \underline{d}_{l_t}^T P_{D^{(t-1)}}^\perp \quad (\text{A.4})$$

$$= \begin{bmatrix} Q_{D^{(t-1)}} & \\ & 0 \end{bmatrix} - \begin{bmatrix} \underline{b}_{t-1} \\ -1 \end{bmatrix} \underline{q}_t^T \quad (\text{A.5})$$

where $\underline{b}_{t-1} = Q_{D^{(t-1)}} \underline{d}_{l_t}$ is the filtered version of \underline{d}_{l_t} using the LS filter $Q_{D^{(t-1)}}$, and $\underline{q}_t^T = \underline{d}_{l_t}^T / \|\underline{d}_{l_t}\|^2$ is the LS filter operator using \underline{d}_{l_t} .

Now, since $\underline{r}_t = P_{D^{(t)}}^\perp \underline{y}_k$ and $\hat{\underline{x}}_k^{(t)} = Q_{D^{(t)}} \underline{y}_k$, premultiplying (A.3) and (A.4) by \underline{y}_k^T yields the recursive update equations for \underline{r}_t and $\hat{\underline{x}}_k(t)$, respectively, i.e.,

$$\underline{r}_t = \underline{r}_{t-1} - \alpha_t \underline{d}_{l_t} \quad (\text{A.6})$$

and

$$\hat{\underline{x}}_k(t) = \begin{bmatrix} \hat{\underline{x}}_k(t-1) \\ 0 \end{bmatrix} + \alpha_t \begin{bmatrix} -\underline{b}_{t-1} \\ 1 \end{bmatrix} \quad (\text{A.7})$$

where $\alpha_t = \underline{q}_t^T \underline{y}_k = \underline{d}_{l_t}^T \underline{r}_{t-1} / \|\underline{d}_{l_t}\|^2$, i.e., the filtered version of \underline{y}_k based upon the LS filter operator \underline{q}_t^T . Thus, the adjustment term in (A.7) corresponding to the previous coefficients is equal to the product of two filtered outputs, namely \underline{b}_{t-1} and α_t , whereas the coefficient associated with the newly added atom is α_t . These equations allow for a ‘‘time-order’’ update after adding a new dictionary atom.

Note that in this algorithm computing the projection matrix $P_{D^{(t)}}$ and LS filter operator $Q_{D^{(t)}}$ are completely avoided and hence no matrix inversion is required. The updated atom is $\underline{d}_{l_t} = P_{D^{(t-1)}}^\perp \underline{d}_{l_t}$; however, we only need the filtered output \underline{b}_{t-1} since

$$\begin{aligned} \underline{d}_{l_t} &= P_{D^{(t-1)}}^\perp \underline{d}_{l_t} \\ &= \underline{d}_{l_t} - D^{(t-1)} Q_{D^{(t-1)}} \underline{d}_{l_t} = \underline{d}_{l_t} - D^{(t-1)} \underline{b}_{t-1}. \end{aligned} \quad (\text{A.8})$$

Once the new basis \underline{d}_{l_t} is selected, steps of the entire algorithm in order of computation are: 1) find $\underline{b}_{t-1} = Q_{D^{(t-1)}} \underline{d}_{l_t}$; 2) compute \underline{d}_{l_t} using (A.8); 3) compute $\underline{q}_t = \underline{d}_{l_t}^T / \|\underline{d}_{l_t}\|^2$; 4) compute α_t ; 5) compute $Q_{D^{(t)}}$ using (A.4); 6) update $\hat{\underline{x}}_k(t)$ using (A.7); and 7) find new \underline{r}_t using (A.6). The order of computations for performing all these equations is $\mathcal{O}(NM)$ as opposed to $\mathcal{O}(N^3)$ for matrix inversion.

ACKNOWLEDGMENTS

The authors would like to thank NSWC-Panama City for support and providing the data used in this study.

REFERENCES

- [1] S. Perry, *Applications of Image Processing to Mine Warfare Sonar*. Melbourne, Australia: Aeronautical and Maritime Research Laboratory, 2000, Ch. 3.
- [2] S. Reed, Y. Petillot, and J. Bell, ‘‘Automated approach to classification of mine-like objects in sidescan sonar using highlight and shadow information,’’ *Proc. Inst. Electr. Eng.—Radar Sonar Navig.*, vol. 151, no. 1, pp. 48–56, Feb. 2004.
- [3] D. Williams, ‘‘Fast target detection in synthetic aperture sonar imagery: A new algorithm and large-scale performance analysis,’’ *IEEE J. Ocean. Eng.*, vol. 40, no. 1, pp. 71–92, Jan. 2015.
- [4] G. Dobeck, J. Hyland, and L. Smedley, ‘‘Automated detection/classification of sea mines in sonar imagery,’’ *Proc. SPIE—Int. Soc. Opt. Eng.*, vol. 3079, pp. 90–11, 1997.
- [5] W. Guo and W. Szymczak, ‘‘Multiresolution neural networks for mine detection in side scan sonar images,’’ *Proc. SPIE—Int. Soc. Opt. Eng.*, vol. 3392, pp. 297–305, 1998.
- [6] J. Tucker, ‘‘Coherence-based underwater target detection for side-scan sonar imagery,’’ M.S. thesis Colorado State Univ., Boulder, CO, USA, May 1999.
- [7] A. Pezeshki, M. R. Azimi-Sadjadi, L. L. Scharf, and M. Robinson, ‘‘Underwater target classification using canonical correlations,’’ in *Proc. OCEANS Conf.*, vol. 4, Sep. 2003, pp. 1906–1911.
- [8] A. Pezeshki, M. R. Azimi-Sadjadi, and L. L. Scharf, ‘‘Undersea target classification using canonical correlation analysis,’’ *IEEE J. Ocean. Eng.*, vol. 32, no. 4, pp. 948–955, Oct. 2007.
- [9] J. Cartmill, N. Wachowski, and M. R. Azimi-Sadjadi, ‘‘Buried underwater object classification using a collaborative multiaspect classifier,’’ *IEEE J. Ocean. Eng.*, vol. 34, no. 1, pp. 32–44, Jan. 2009.
- [10] J. Tucker and M. Azimi-Sadjadi, ‘‘Coherence-based underwater target detection from multiple disparate sonar platforms,’’ *IEEE J. Ocean. Eng.*, vol. 34, pp. 38–52, 2011.
- [11] N. Klausner and M. Azimi-Sadjadi, ‘‘Detection in multiple disparate systems using multi-channel coherence analysis,’’ *IEEE Trans. Aerosp. Electron. Syst.*, vol. 48, pp. 3554–3566, Oct. 2012.
- [12] L. Scharf and C. Mullis, ‘‘Canonical coordinates and the geometry of inference, rate, and capacity,’’ *IEEE Trans. Signal Process.*, vol. 48, pp. 824–891, Mar. 2000.
- [13] D. Cochran, H. Gish, and D. Sinno, ‘‘A geometric approach to multiple-channel signal detection,’’ *IEEE Trans. Signal Process.*, vol. 43, pp. 2049–2057, Sep. 1995.
- [14] A. Neilsen, ‘‘Multiset canonical correlations analysis and multispectral, truly multitemporal remote sensing data,’’ *IEEE Trans. Image Process.*, vol. 11, p. 293–305, Mar. 2002.
- [15] N. Klausner, M. Azimi-Sadjadi, and L. Scharf, ‘‘Detection of spatially correlated time series from a network of sensor arrays,’’ *IEEE Trans. Signal Process.*, vol. 62, pp. 1396–1407, Jan. 2014.
- [16] N. Klausner, ‘‘Detection of multiple correlated time series and its application in synthetic aperture sonar imagery,’’ Ph.D. dissertation, Colorado State Univ., Fort Collins, CO, USA, May 2014.
- [17] L. Scharf and B. Friedlander, ‘‘Matched subspace detectors,’’ *IEEE Trans. Signal Process.*, vol. 42, no. 8, pp. 2146–2157, Aug. 1994.
- [18] M. Aharon, M. Elad, and A. Bruckstein, ‘‘K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation,’’ *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, 2006.
- [19] S. S. Chen, D. D. L. Donoho, and M. A. Saunders, ‘‘Atomic decomposition by basis pursuit,’’ *SIAM Rev.*, vol. 43, no. 1, pp. 129–159, 2001.
- [20] S. Mallat and Z. Zhang, ‘‘Matching pursuits with time-frequency dictionaries,’’ *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [21] J. Alder, B. Rao, and K. Kreutz-Degado, ‘‘Comparison of basis selection methods,’’ in *Conf. Record 13th Asilomar Conf. Signal Syst. Comput.*, 1996, vol. 1, pp. 252–257.
- [22] A. M. Bruckstein, D. Donoho, and M. Elad, ‘‘From sparse solutions of systems of equations to sparse modelling of signals and images,’’ *SIAM Rev.*, vol. 51, no. 1, pp. 34–81, 2009.
- [23] A. Bruckstein, D. Donoho, and M. Elad, ‘‘From sparse solutions of systems of equations to sparse modeling of signals and images,’’ *SIAM Rev.*, vol. 51, no. 1, pp. 34–81, 2009.
- [24] Q. Zhang and B. Li, ‘‘Discriminative K-SVD for dictionary learning in face recognition,’’ in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2010, pp. 2691–2698.
- [25] M. R. Azimi-Sadjadi, J. Kopacz, and N. K. Klausner, ‘‘K-SVD dictionary learning using a fast OMP with applications,’’ in *Proc. IEEE Int. Conf. Image Process.*, Oct. 2014, pp. 1599–1603.
- [26] I. Markovsky, *Low-Rank Approximation: Algorithms, Implementation, and Applications*. London, U.K.: Springer-Verlag, 2011, Ch. 2.4.
- [27] D. Brown, D. Cook, and J. Fernandez, ‘‘Results from a small synthetic aperture sonar,’’ in *Proc. OCEANS Conf.*, 2006, DOI: [10.1109/OCEANS.2006.306855](https://doi.org/10.1109/OCEANS.2006.306855).
- [28] L. Scharf, *Statistical Signal Processing: Detection, Estimation, and Image Series Analysis*. New York, NY, USA: Addison-Wesley, 1991, Ch. 2.6.

- [29] J. Kerekes, "Receiver operating characteristic curve confidence intervals and regions," *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 2, pp. 251–255, 2008.
- [30] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," 1996, pp. 226–231
- [31] S. T. Alexander, *Adaptive Signal Processing: Theory and Applications*. New York, NY, USA: Springer-Verlag, 1984, Ch. 8.



Mahmood R. Azimi-Sadjadi (M'81–SM'89) received the M.S. and Ph.D. degrees in electrical engineering with specialization in digital signal/image processing from the Imperial College of Science and Technology, University of London, London, U.K., in 1978 and 1982, respectively.

He is currently a Full Professor with the Electrical and Computer Engineering Department, Colorado State University (CSU), Fort Collins, CO, USA. He is also serving as the Director of the Digital Signal/Image Laboratory at CSU. His main areas of in-

terest include statistical signal and image processing, target detection, classification and tracking, sensor array processing, learning systems, and distributed sensor networks. He is the c-author of the book *Digital Filtering in One and Two Dimensions* (New York, NY, USA: Plenum Press, 1989).

Dr. Azimi-Sadjadi served as an Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING and the IEEE TRANSACTIONS ON NEURAL NETWORKS.



Nick Klausner (S'08–M'14) received the M.S. and Ph.D. degrees in electrical and computer engineering from Colorado State University, Fort Collins, CO, USA, in 2010 and 2014, respectively.

Since 2014, he has been a Research Scientist at Information System Technologies, Inc., Fort Collins, CO, USA. His current research interests include statistical signal and image processing.



Justin Kopacz received the B.S. and M.S. degrees in electrical and computer engineering from Colorado State University, Fort Collins, CO, USA, in 2013 and 2014, respectively.

His research interests include statistical signal and image processing.