

Symbolic Analysis of Sonar Data for Underwater Target Detection

Kushal Mukherjee, *Student Member, IEEE*, Shalabh Gupta, *Member, IEEE*, Asok Ray, *Fellow, IEEE*, and Shashi Phoha, *Senior Member, IEEE*

Abstract—This paper presents a symbolic pattern analysis method for robust feature extraction from sidescan sonar images that are generated from autonomous underwater vehicles (AUVs). The proposed data-driven algorithm, built upon the concepts of *symbolic dynamics* and *automata theory*, is used for detection of mines and mine-like objects in the undersea environment. This real-time algorithm is based on symbolization of the data space via coarse graining, i.e., partitioning of the two-dimensional sonar images. The statistical information, in terms of stochastic matrices that serve as features, is extracted from the symbolized images by construction of probabilistic finite state automata. A binary classifier is designed for discrimination of detected objects into mine-like and nonmine-like categories. The pattern analysis algorithm has been validated on sonar images generated in the exploration phase of a mine hunting operation; these data have been provided by the Naval Surface Warfare Center. The algorithm is formulated for real-time execution on limited-memory commercial-of-the-shelf platforms and is capable of detecting objects on the seabed-bottom.

Index Terms—Feature extraction, mine countermeasures, pattern classification, sonar data analysis, symbolic dynamics, target detection.

I. INTRODUCTION

RAPID advancement of modern engineering technology has led to development of increasingly portable manned and unmanned platforms for mine countermeasure (*MCM*) operations of the naval forces. In particular, autonomous underwater vehicles (*AUV*) provide enhanced speed and improved search efficiency in *MCM* operations. These vehicles are equipped with advanced sensing devices, including sidescan sonar imaging systems, to search a target area for mines lying on the seabed. Apart from mines, several objects that are commonly found on the seabed may also have similar shape and structure as those of mines. In this paper, these objects are collectively referred to as mine-like objects.

Manuscript received August 11, 2010; revised December 20, 2010; accepted February 23, 2011. Date of publication May 12, 2011; date of current version May 27, 2011. This work was supported in part by the U.S. Office of Naval Research under Grant N00014-09-1-0688, and by the U.S. Army Research Laboratory and U.S. Army Research Office (ARO) under Grant W911NF-07-1-0376. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the sponsoring agencies.

Associate Editor: N. Chotiros.

The authors are with the Pennsylvania State University, University Park, PA 16802 USA (e-mail: kum162@psu.edu; szg107@psu.edu; axr2@psu.edu; sxp26@psu.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JOE.2011.2122590

The current state-of-the-art *MCM* techniques involve deployment of AUVs that perform sequential operations of detection, classification, marking, and neutralization [1]. Sidescan sonar systems are used for efficient mapping of large areas of the sea floor to generate sonar imagery [2] that contains information about the seabed-bottom, typical objects of interest (e.g., undersea mines), and the clutter. Here the term “clutter” refers to general seabed debris (e.g., rocks, wreckage, and sand ripples). Traditionally, sonar imagery generated by sidescan sonars mounted on the AUVs are analyzed offline in batches before the next operation (e.g., classification). In general, the exploration phase of mine hunting involves scanning vast areas of the sea bed to detect very sporadic occurrences of mines or mine-like objects. To expedite mine-hunting operations in the exploration phase, AUVs perform long-range scanning with low-resolution sonars. Therefore, very limited resolution data are available in the exploration phase for detection of mines or mine-like objects. Once such an object is detected, a higher resolution short-range sonar is used to take a second look at the suspected object [3]. Thus, upon preliminary detection of an object, the next phase of exploitation in a mine hunting operation generates scan data of much higher resolution to ascertain whether the detected object is truly a mine or not and to which class of mines it belongs. Commercially available synthetic aperture sonar (SAS) [4], [5] is a potential candidate for generating such high resolution imagery.

This paper addresses the problem of detection of mine-like objects from a low-resolution sonar images. The example problem specifically addresses the situation in the exploration phase where the available data is very limited. However, the underlying theory of mine detection, presented in this paper, can be extended to handle high-resolution scan data for accurate mine detection and classification; this is a topic of future research.

Mine-like objects are often detected in sonar images by using the highlight and the shadow cast by the object. The shadow is usually a critical feature for distinguishing mine-like objects from the background. In the traditional approach to mine detection, an adaptive threshold is often assigned to the mapped features based on the fact that the specular reflection from the target is brighter (i.e., a stronger reflected signal) than the diffused reflection from the seabed. This is the simplest form of non-parametric detection and has been shown to be effective for object detection [6] in a relatively clutter-free environment. However, a cluttered background may generate many false-positive mine detections. Recent mine detection methods have made use of advanced signal processing techniques. For example, Reed

et al. [7] have used Dempster–Shafer information to classify mines for reduction of false alarms; and Dura *et al.* [8] have proposed a data-adaptive algorithm to eliminate the need for *a priori* training of MCM missions. A feature level fusion method for detection of land mines has been reported in literature, which uses a Bayesian framework involving likelihood ratio tests [9]. Improvements to traditional Bayesian detection methods, including usage of geometric and statistical properties of objects have been proposed by Calder *et al.* [10]. Statistical modeling provides considerable speckle filtering while retaining useful information [11]. An image-classification technique, proposed by Bhradwaj and Carin [12], employs templates tailored to image subcomponents; a Markov model has been used to statistically characterize the relationship between neighboring image subcomponents. In another paper, advanced image classification tools have been developed based on hidden conditional random fields [13]. This method employs a neighborhood graph structure to capture conditional dependencies between regions of the image. Ye *et al.* [14] have used level-set models to segment sonar images into the shadow and highlight regions. This method is based on features extracted by using Gaussian Markov random fields and effectively removes speckle noise from the sonar image. Maussang *et al.* [15], [16] have applied adaptive data thresholding for object detection as well as statistical methods that do not require the presence of a shadow for mine detection, which is useful for detection of buried mines. Dee *et al.* [17] have presented a survey of the recent advances in the application of graphical models (e.g., Markov Bayesian networks) for the purpose of real-time visual surveillance.

Recently the authors and their colleagues have developed a real-time data-driven pattern identification method, called *symbolic pattern analysis* [18]–[20], which enables information compression from large data sets into statistical patterns of low dimensions. The underlying principles of this method are built upon multidisciplinary concepts derived from *symbolic dynamics* [21], *statistical mechanics* [20], [22], *pattern recognition* [23], *information theory* [24], and *automata theory* [25]. The algorithms of symbolic pattern analysis method have been experimentally validated for real-time execution in diverse applications, including electronic circuits [19], fatigue damage monitoring in polycrystalline alloys [26], [27], and robot signature analysis [28]. This technique has been shown to perform very well for pattern identification as compared to Bayesian analysis, neural networks, principle component analysis, and other statistical methods in terms of speed of execution, memory requirements, sensitivity to small changes in relevant information, and robust detection in the presence of noise [29]. In these applications, the symbolic pattern analysis algorithms have been applied to one-dimensional sensor signals.

The primary advantages of symbolic pattern analysis method are robustness to noise and fast computation on commercially available inexpensive platforms that are suitable for real-time applications. This method facilitates onboard implementation of the underlying algorithms for in situ pattern analysis, thereby enhancing autonomy of underwater vehicles and saving the time of mission operations.

The contributions of the current paper beyond the previously reported work (i.e., symbolic analysis of one-dimensional time series data [18]–[20]) are listed below:

- extension of the concepts of symbolic pattern analysis for two-dimensional (i.e., image) data analysis for pattern identification;
- application and validation of the proposed method on a set of sonar data in the form of low-resolution images for underwater mine detection.

The paper is organized as follows. Section II presents a geometric model of mines for pattern analysis. Section III presents the concepts of symbolic pattern analysis including brief reviews of symbolic dynamics, space partitioning methods, and construction of probabilistic finite state automata (PFSA) for development of the mine detection algorithm. This section also presents the methods of feature extraction and classifier construction using the symbolic pattern analysis algorithm. Section IV presents the results of mine detection in sonar images based on a set of test data provided by Naval Surface Warfare Center (NSWC), Panama City, FL. Section V concludes the paper along with recommendations for future work.

II. GEOMETRIC MODEL FOR MINE DETECTION

This section presents a geometric model for detection of underwater mines from the sidescan sonar data sets in the form of images, which are provided by NSWC. Each sonar image has a size of 512×1000 pixels (512 pixels along the range of the sidescan sonar and 1000 pixels along the direction of motion of the AUV). The range of the sonar system, used in this paper, is approximately 30 m. A single pixel represents an approximate rectangular area of $6 \text{ cm} \times 3 \text{ cm}$. Note that the true dimension of the pixels becomes wider with an increase in the range of the sonar beam; it is assumed that small variations in the range have no significant effects on the performance of the algorithm. The inclusion of these effects in the data-analysis algorithm is considered as a future work in Section V.

The ground-truth information about the exact location of mines is available for a set of 151 images that were acquired by an AUV. The ensemble of data sets is divided into a training set of 91 images (consisting of in total 100 mines) and a test set of 60 images for validation. The parameters required for the symbolic pattern analysis algorithm, and the receiver operating characteristics (ROC) curve are obtained based on the ground-truth statistics provided in the training data set (see Section III).

Fig. 1 shows the operation of a typical sidescan sonar. The sonar moves along a straight track with a constant speed and at a constant height above the seabed. Only the left half of the transmitted sonar beam is shown here in Fig. 1 and the right half is similar. The transducers on both sides transmit narrow beams of energy approximately perpendicular to the direction of platform's motion. Raw sidescan imagery corresponds to acoustic echo intensity versus time of flight. The horizontal range is deduced from the raw imagery by assuming that the seabed is horizontal (i.e., zero gradient). Fig. 2 shows the sonar ray propagation model along with a typical image of a mine present on the seabed. The echo from the mine-like object (present in the

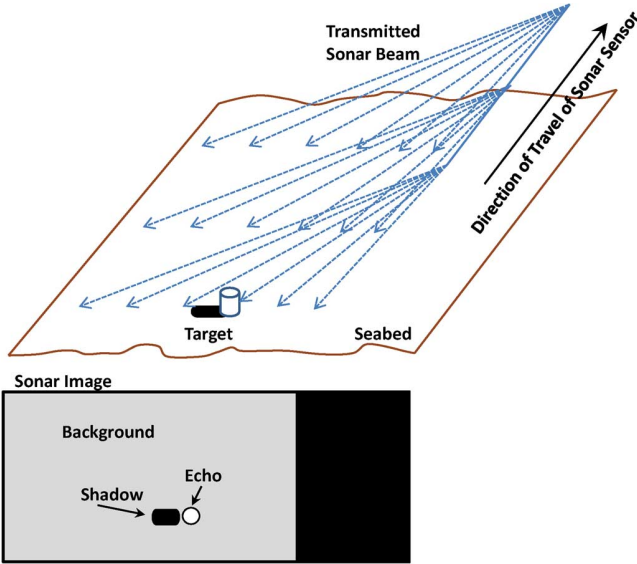


Fig. 1. Illustration of the operation of a typical sidescan sonar. The transmitted sonar beam on the left side of the platform is shown here. A similar sonar beam on the right side of the platform is not shown. Reflections of the sonar beam from the objects on the seabed are also not shown in this figure for clarity of presentation.




Fig. 2. Shadow formation due to a mine-like object in the path of sonar waves. The echo from the mine-like object is mapped to the range marked A' to B' .

range marked A to B) is mapped to the range marked A' to B' because of the assumption of a horizontal seabed.

A geometric model has been determined for feature extraction to detect mines in a sonar image. This model is used in the training set to obtain the distributions of the various features that characterize a typical mine and in the testing set as a template for detecting the mines. Based on the principles of sidescan sonar technology and properties of sound wave propagation in the ocean, a mine is characterized by two distinct regions that are adjacent to each other in the sonar image [10]. These regions correspond to a highlight (i.e., an echo) and a shadow. It is assumed that there are no refraction effects below the surface of water. This assumption is valid for relatively short ranges (~ 30 – 50 m) and high-frequency acoustic signals, provided that multipath returns from surface reflections are small; in this case, the sonar rays can be approximated to propagate in straight lines. The principle of shadow formation when an obstacle falls in the path of sidescan sonar rays is illustrated in Fig. 2. The region, adjacent to the mine and away from the sonar falls within the shadow cast by the mine and appears dark.

The reflections from a mine-like object are usually much stronger than those from reverberation or those caused by the clutter. As described above, there is a sonar shadow because of the structure of the mine protruding off the seabed. This shadow is a very good indicator for distinguishing mines from background features that include seabed and surface reverberation and clutter. A mine usually appears as a highlight; and since the mines under consideration are placed on the seabed, there is a shadow region that exists adjacent to a mine in the direction away from the sidescan sonar receiver. The regions adjacent to the mine and the shadow in the direction perpendicular to the sonar rays on the seabed, are also important features. These regions help in distinguishing mines from reverberation and other objects (e.g., large rocky structures on the seabed) that may generate highlights and also possibly cast shadows behind them. Therefore, for improving the detection performance, these background regions that are adjacent to the mine and the shadow are divided into two regions, namely, the mine background and the shadow background. Thus, the geometric model for mine detection consists of the following four distinct regions:

- 1) *main body of the mine*- denoted as A^m ;
- 2) *shadow region near the mine*- denoted as A^s ;
- 3) *mine Background*- denoted as A^{mb} ;
- 4) *shadow Background*- denoted as A^{sb} .

If an isolated mine is present, then the region A^m is characterized by a highlight and the shadow region A^s by dark pixels. The shadow length depends on the range of the mine from the sensor. However, in the current set of images, the variations in the shadow length are not large. Consequently, the length of A^s is assumed to be fixed and is based on the assumption of a nominal mine being present at mid range. The statistical pattern analysis tool, introduced in this paper, is robust to variations in the shadow length. Construction of flexible geometric models is a topic of future research.

The background regions A^{mb} and A^{sb} are characterized by moderate intensity pixels. Thus, for differentiating mine-like objects with large size objects (e.g., ship wreckage or rocky bottoms), the background regions adjacent to the mine and the shadow must show different characteristics than those of the mine and the shadow, respectively. Further details are presented in Section III-D.

Fig. 3 displays the geometric model with the number of pixels allocated to each region of the model. It is assumed that mines are spatially isolated from each other and that they are surrounded by the general seabed background. The model for mine detection is parameterized based on the following aspects: 1) geometric properties of anticipated objects; 2) size parameters estimated from the objects present in the ground truth; and 3) physical understanding of the relative importance of the various components of the geometric model. Model parametrization reduces the mine detection process to comparison of various statistics generated from the ground-truth analysis.

III. SYMBOLIC PATTERN ANALYSIS

This section presents the underlying concepts and salient properties of the *symbolic pattern analysis* method. While the details of this method have been reported in previous

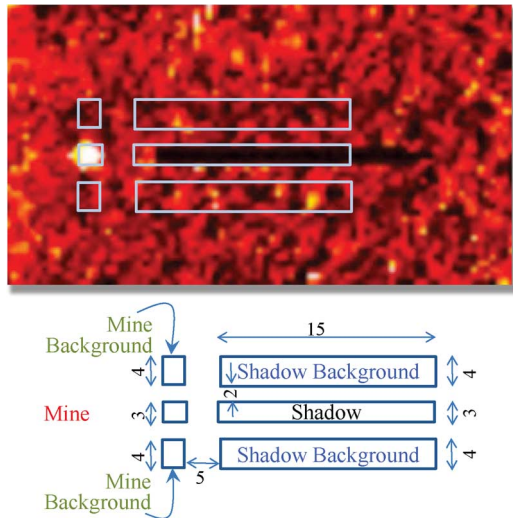


Fig. 3. Geometric model for mine detection from sonar images.

publications for one dimensional time series data analysis [18], [20], this paper extends the concepts of symbolic pattern analysis for two-dimensional data (i.e., an image) and presents an application for detection of undersea mines in sonar images. The major steps of symbolic pattern analysis are as follows:

- encoding nonlinear system dynamics of the underlying process by partitioning the observed time series data for generation of symbol sequences;
- construction of a probabilistic finite state automata (PFSA) from the symbol sequence;
- generation of pattern vectors as features for detection and classification tasks.

A brief review of symbolic pattern analysis method for feature extraction from time series data is presented below.

In the symbolic dynamics literature [18], [21], the primary objective is to represent time series data from a dynamical system as a symbol sequence. Let $\Omega \in \mathbb{R}^n$ be a compact (i.e., closed and bounded) region, within which the trajectory of the dynamical system is circumscribed. The region Ω is partitioned into a finite number of (mutually exclusive and exhaustive) cells, so as to obtain a coordinate grid. Let the cell, visited by the trajectory, be denoted as a random variable taking a symbol value from the alphabet Σ . An orbit of the dynamical system is described as $\{x_0, x_1 \cdots x_k \cdots\}$ with $x_i \in \Omega$, which passes through or touches one of the cells of the partition. The symbol sequence is denoted as $\{\sigma_0, \sigma_1, \cdots \sigma_k \cdots\}$ where each symbol σ_i belongs to the (finite) alphabet set Σ . Symbolic dynamics may be viewed as coarse graining of the data space, which is subjected to (possible) loss of information resulting from granular imprecision of partitioning boxes. However, the essential robust features are preserved in the symbol sequences through an appropriate partitioning of the region Ω .

A crucial step in analysis of sensor signals based on the concepts of symbolic dynamics is to partition the time series data sets for symbol sequence generation. Several partitioning techniques have been reported in literature for symbol generation. These techniques are primarily based on symbolic false nearest neighbors (*SFNN*) [30], which may become cumbersome and

extremely computation-intensive if the data set is contaminated by noise. As an alternative, the time series data can be partitioned for symbol sequence generation using either the uniform or the maximum-entropy partitioning [19]. For certain applications, prior to partitioning the time series data may be processed by using an appropriately chosen wavelet transform [26].

Once the symbol sequence is generated, the next step is construction of a probabilistic finite state automata for modeling the statistical dependencies between the occurrences of symbols. The symbolic sequence is modeled as a D -Markov process where the likelihood of a future symbol depends only on the previous D symbols [18]. This assumption leads to construction of probabilistic finite state automata (PFSA), wherein the state transition probabilities depict the conditional dependencies between symbols. The state probability vectors, which are derived from the state transition probability matrices of the PFSA, serve as statistical patterns of the underlying process. The algorithms of the symbolic pattern analysis method have been experimentally validated for real-time execution in diverse applications, including electronic circuits [18], fatigue damage monitoring in polycrystalline alloys [26], [27] and robot signature analysis [28].

This paper aims to extend the concepts of symbolic pattern analysis for the purpose of detection of underwater mines in sonar images; the fundamental difference being that the image data is two dimensional unlike time series data that has a one dimensional domain, i.e., time. The essential concepts of space partitioning for image symbolization, construction of probabilistic finite state automata (PFSA), and extraction of the features are succinctly described below.

A. Construction of a Physics-Based Partitioning

This paper uses a physics-based partitioning scheme that associates a relevant meaning to each cell of the partitioning. Two important parameters need to be determined for successful application of symbolic dynamics to address the problem of mine detection. The first parameter is the alphabet size $|\Sigma|$ and the second is the vector parameter of the partition segment boundaries. In the particular application of mine detection in sonar images, the essential robust features that need to be preserved are: the highlights, i.e., the reflections from the front of an object protruding above the seabed, and the long shadow that follows the object. Apart from these, the average background feature generated from the general seabed reverberation and the clutter has also been used in this paper for pattern analysis.

The three plates in Fig. 4 shows the histograms of frequency counts of pixel intensities for mine, shadow, and the combined background regions adjacent to the mine and the shadow in the geometric model (see Section II). The histograms depict the number of pixels versus pixel intensity (0 to 255) for each region of the geometric model. The observed distribution of pixel intensities in the background regions of the mine and the shadow are similar; therefore, they have been combined in a single histogram of the background region, as shown in the right hand plate of Fig. 4; the background region is largely dominated by medium-intensity pixels. These histograms are generated from the *a priori* known ground-truth information about the location of mines from the training data set of 91 sonar images. The mine

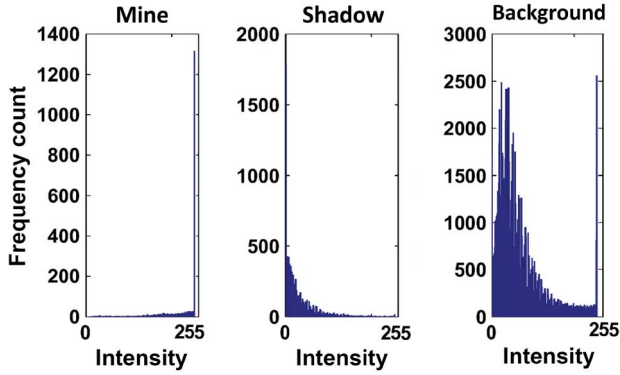


Fig. 4. Histograms of sonar wave reflections: 1) the left plate represents mines; 2) the middle plate represents shadows; and 3) the right plate represents the background including the reverberation and the clutter. The distribution of pixel intensities in the regions of mine background and shadow background are combined together in a single histogram of the background region, because they show similar trends.

histogram in the left hand plate of Fig. 4 shows that the pixels in the mine regions generally have a high intensity. On the other hand, the shadow histogram in the middle plate of Fig. 4 depicts that the pixels in the shadow region have a low intensity. The histograms for the three regions are reasonably separated, although some regions of the background histogram overlap with the mine and the shadow histograms, which is not of significant concern because the three regions in the geometric model are separated in space. Therefore, for partitioning the sonar images, a distinct symbol in the alphabet $\Sigma = \{a, b, c\}$ is assigned to each of the three features corresponding to mine, shadow, and the background. In this application, an alphabet of three symbols, i.e., $|\Sigma| = 3$, is seen to be adequate for characterizing the above features because the information gained by having $|\Sigma| > 3$ is found to be insignificant to offset the additional computation. Furthermore, an increased $|\Sigma|$ would require a larger amount of data to achieve similar accuracy in estimation of the stochastic matrix, described later in Section III-B. As the complexity of the background histogram increases, it may be necessary to increase the number of symbols.

The next important consideration is selection of the segment boundaries for partitioning the data sets. Traditional partitioning techniques (e.g., uniform partitioning and maximum entropy partitioning [19]) may not adequately capture the details of mine patterns. Therefore, this paper has adopted a different partitioning method that makes use of the physics-based information from the statistics as depicted in the three histograms of Fig. 4. Based on the information contained in these histograms, the partitioning is constructed by assigning a symbol a to high-intensity pixels ranging from 171 to 255 on the gray scale; similarly, symbol b is assigned to the medium intensity pixels ranging from 51 to 170; and symbol c is assigned to low intensity pixels ranging from 0 to 50. As seen from Fig. 4, approximately 98% of the pixels in the histogram, generated from the mine region in the geometric model, correspond to the symbol a (i.e., bright pixels). Similarly, approximately 80% of the pixels in the histogram, generated from the shadow region in the geometric model, correspond to the symbol c (i.e., dark pixels). A majority of the pixels in the background regions of

the geometric model correspond to the symbol b (i.e., moderately dark pixels). Thus, the entire image is symbolized and represented by a two-dimensional array of symbols belonging to the alphabet $\Sigma = \{a, b, c\}$. Symbolization of an image is described by the following definitions.

Definition 1: Let $\mathcal{H} \triangleq \{(i, j) : i, j \in \mathbb{N}_1 \triangleq \{1, 2, \dots\} \mid 1 \leq i \leq m, 1 \leq j \leq n\}$ be the set of coordinates of an image of size $m \times n$. Let $A \triangleq \{0, 1, \dots, 255\}$ be the set of gray-scale intensity levels. Then, an image \mathcal{I} is defined by a map $\mathcal{I} : \mathcal{H} \mapsto A$.

Definition 2: Let Σ be the partitioning alphabet and partitioning of the set A be defined by a map $P : A \mapsto \Sigma$ such that each pixel value of the image is mapped to a symbol in Σ . Then, the symbolized image is defined by a map $\mathcal{I}_\Sigma \equiv P \circ \mathcal{I}$ such that $\mathcal{I}_\Sigma : \mathcal{H} \mapsto \Sigma$. For a nonempty region $\mathcal{B} \subseteq \mathcal{H}$, its configuration is denoted as $\mathcal{I}_\Sigma(\mathcal{B})$.

The above partitioning scheme enables robust detection of mines with a high probability of detection and a very low probability of false alarms as discussed later in Section IV. Furthermore, this method of symbolization significantly reduces the memory requirements. The next subsection describes the construction of a probabilistic finite state automata (PFSA) for feature extraction using the geometric model of mine detection.

B. Neighborhood Modeling

In the symbolized image \mathcal{I}_Σ generated from a typical sonar data set, every gray-scale pixel is labeled by a symbol $\sigma \in \Sigma = \{a, b, c\}$. The symbols a , b and c represent three different intensity levels on the gray scale. Let $\mathcal{B} \subset \mathcal{H}$ represent a local region, i.e., one of the four regions: the mine (A^m), the shadow (A^s), the mine background (A^{mb}), and the shadow background (A^{sb}). For each of these regions of the geometric model, a probabilistic finite state automaton (PFSA) is constructed for local feature extraction from sonar images in terms of the neighborhood of a pixel $(i, j) \in \mathcal{B}$. For example, a neighborhood may be constructed by including all the elements that form a 3×3 square centered at the pixel (i, j) .

Definition 3: In the discrete topology, the κ -neighborhood of a pixel at a location $(i, j) \in \mathcal{B}$ is defined as

$$\mathcal{N}_\kappa(i, j) = \{(i', j') \in \mathcal{H} : \max(|i - i'|, |j - j'|) \leq \kappa\} \quad (1)$$

where $\kappa \in \mathbb{N}_0 \triangleq \{0, 1, 2, \dots\}$ is the neighborhood radius. Note: $\mathcal{N}_0(i, j)$ is the singleton set $\{(i, j)\}$.

Definition 4: Let σ_{ij} be the symbol at a pixel location (i, j) , i.e., $\sigma_{ij} = \mathcal{I}_\Sigma(i, j)$, for an image \mathcal{H} . Then, the configuration $\mathcal{I}_\Sigma(\mathcal{H})$ is modeled as a Markov random field (MRF) if the following two conditions on the probability measure $P(\bullet)$ hold [31]:

- 1) $P(\mathcal{I}_\Sigma(\mathcal{H})) > 0$;
- 2) $P(\sigma_{ij} | \mathcal{I}_\Sigma(\mathcal{H} / \{(i, j)\})) = P(\sigma_{ij} | \mathcal{I}_\Sigma(\mathcal{N}_\kappa(i, j) / \{(i, j)\})) \quad \forall (i, j) \in \mathcal{H}$.

Remark 1: The MRF properties ensure that the probability of a symbol depends only on the configuration of its κ -neighborhood, thereby eliminating long-range dependencies.

Definition 5: For $\kappa \in \mathbb{N}_1$, let $\mathcal{N}_\kappa(i, j)$ be a κ -neighborhood of a pixel location $(i, j) \in \mathcal{H}$. Then, the (Markov) state q_{ij} is represented by the configuration $\mathcal{I}_\Sigma(\mathcal{N}_{(\kappa-1)}(i, j))$.




Fig. 5. Conversion from symbols (in the alphabet Σ) to compressed states (in the state set S^c). Note: The center state is labeled as s_1 following the state compression methodology.

The above definition of state implies that the translation of a $(\kappa - 1)$ -neighborhood window covers the κ -neighborhood. Let the set of all possible configurations in a $(\kappa - 1)$ -neighborhood be denoted as the set $\mathcal{S} \triangleq \{s_1, s_2 \dots s_{|\mathcal{S}|}\}$ of states, where $|\mathcal{S}|$ is cardinality of \mathcal{S} , i.e., the total number of states. Thus, at every point $(i, j) \in \mathcal{H}$, the $(\kappa - 1)$ -neighborhood window corresponds to a particular state $q_{ij} \in \mathcal{S}$ on the image, and $|\mathcal{S}|$ is bounded above by $|\Sigma|^{|\mathcal{N}_{(\kappa-1)}|}$, i.e., $|\mathcal{S}| \leq |\Sigma|^{|\mathcal{N}_{(\kappa-1)}|}$; the inequality is because of the fact that some of the states might have zero probability of occurrence.

In this process, every pixel in a region \mathcal{H} can be mapped to a state, provided that the pixel is not at the periphery of the image. The transformation of neighborhood symbols to a state is illustrated in Fig. 5, where the state is chosen as the configuration of $\mathcal{N}_{(\kappa-1)}$ with the underlying assumption of a Markov random field (see Definition 4). Therefore, the “state image” has effective interactions only between the nearest neighbors.

In general, a large number of states require a high computational capability and hence might not be feasible for real-time applications. Both the size (in pixels) of a state and cardinality $|\mathcal{S}|$ of the state set increase with the radius κ of the neighborhood and cardinality $|\Sigma|$ of the alphabet. For example, if $\kappa = 2$ and $|\Sigma| = 3$, then a state consists of $|\mathcal{N}_{(\kappa-1)}| = 9$ (i.e., 3×3) pixels and $|\mathcal{S}| \leq |\Sigma|^{|\mathcal{N}_{(\kappa-1)}|} = 3^9$ pixels. Therefore, for computational efficiency, it is necessary to compress the set of states \mathcal{S} to an effective smaller set S^c by state merging [21] that enables mapping of two or more different configurations to a single state. A state merging must preserve enough information as needed for mine detection, even though it may lead to lossy coding of the image. State equivalence classes are now defined in the context of state merging.

Let (i, j) and (i', j') be two different locations on an image. The configurations of symbols in $\mathcal{N}_{(\kappa-1)}(i, j)$ and $\mathcal{N}_{(\kappa-1)}(i', j')$ are represented by the same *state* if they have identical occurrences of each symbol $\sigma \in \Sigma$. In other words, the state denoted by the configuration of symbols in a neighborhood depends only on the types and the number of symbols and not on the arrangement of symbols in the $(\kappa - 1)$ -neighborhood. In the previous example, for $|\Sigma| = 3$ and $\kappa = 2$, the operation of state merging reduces the number of states to $|S^c| = (|\mathcal{N}_{(\kappa-1)}| + |\Sigma| - 1)! / (|\mathcal{N}_{(\kappa-1)}|!(|\Sigma| - 1)!) = 55$ as compared to $|\mathcal{S}| = 3^9$ without state merging. A formal definition of state merging follows.




Fig. 6. Example of a 3-state probabilistic finite state automaton (PFSA).

Definition 6: For the set \mathcal{S} of states, any two states $s_i \in \mathcal{S}$ and $s_j \in \mathcal{S}$ belong to the same equivalence class \mathcal{E} if

$$T_{s_i}(\sigma) = T_{s_j}(\sigma) \quad \forall \sigma \in \Sigma \quad (2)$$

where $T_{\bullet}(\star)$ denotes the number of occurrences of symbol $\star \in \Sigma$ in a state $\bullet \in \mathcal{S}$.

Let the number of states after state compression be denoted by S^c . Subsequently, a probabilistic finite state automaton (PFSA) is constructed from the set S^c such that the states of the PFSA are the elements of the compressed set S^c and the edges are the elements of the $|S^c| \times |S^c|$ left stochastic matrix $\Pi^{\mathcal{B}} \equiv [\pi_{jk}^{\mathcal{B}}]$, where $\pi_{jk}^{\mathcal{B}} = P(s_j|s_k)|_{\mathcal{B}}$, i.e., the probability of the state s_j being a neighbor of the state s_k . (The superscript denotes that the state transition matrix $\Pi^{\mathcal{B}}$ is evaluated from the region \mathcal{B} .) [Note: $\pi_{jk}^{\mathcal{B}} \geq 0 \forall j, k$ and $\sum_k \pi_{jk}^{\mathcal{B}} = 1 \forall j$.] The stochastic matrix $\Pi^{\mathcal{B}}$ is given as

$$\Pi^{\mathcal{B}} = \begin{bmatrix} P(s_1|s_1) & \dots & P(s_{|S^c|}|s_1) \\ \vdots & \ddots & \vdots \\ P(s_1|s_{|S^c|}) & \dots & P(s_{|S^c|}|s_{|S^c|}) \end{bmatrix}_{\mathcal{B}} \quad (3)$$

Fig. 6 shows the PFSA for a typical local region $\mathcal{B} \subset \mathcal{H}$. The transition probabilities $P(s_l|s_k)$ are computed numerically as given below

$$P(s_l|s_k) = \frac{N(s_l, s_k)}{\sum_{l'=1,2,\dots,|S^c|} N(s_{l'}, s_k)} \quad \forall s_l, s_k \in S^c \quad (4)$$

where $N(s_l, s_k)$ is the total number of occurrences of the neighboring pairs consisting of states s_l and s_k . Note: $P(s_l|s_k) \neq P(s_k|s_l)$ because $P(s_l|s_k)$ is normalized over s_l while $P(s_k|s_l)$ is normalized over s_k . As a result, the matrix Π is not symmetric in general, however, it is a stochastic matrix.

C. Feature Extraction

The stochastic matrix $\Pi^{\mathcal{B}}$ associated with the PFSA captures the neighborhood interactions between the symbols in a given region \mathcal{B} . This stochastic matrix is treated as a feature based on which the subsequent task of pattern classification is performed.

Construction of PFSA and the corresponding stochastic matrix Π follow the principle of sliding block code [21]. Thus, PFSA are constructed by sliding the geometric model (see




Fig. 7. Concept of feature extraction and classification for mine detection via symbolic pattern analysis.

Section II) on the sonar image as depicted in Fig. 7. For every pixel location (i, j) , the geometric model is constructed around that pixel such that (i, j) lies at the center point of the mine region. For a certain position of the geometric model on the image, a PFSA is constructed for each region of the model. For any pixel location (i, j) on the sonar image, the following four features (stochastic matrices) are obtained:

- 1) $\Pi^m(i, j)$ is constructed from the mine region (A^m) where pixel location (i, j) lies at the center of the mine region (A^m) in the geometric model;
- 2) $\Pi^s(i, j)$ is constructed from the shadow region (A^s) where pixel location (i, j) lies at the center of the mine region (A^m) in the geometric model;
- 3) $\Pi^{mb}(i, j)$ is constructed from the background region adjacent to the mine (A^{mb}), where pixel location (i, j) lies at the center of the mine region (A^m) in the geometric model;
- 4) $\Pi^{sb}(i, j)$ is constructed from the background region adjacent to the shadow (A^{sb}), where pixel location (i, j) lies at the center of the mine region (A^m) in the geometric model.

Thus, each pixel location (i, j) is associated with four stochastic matrices corresponding to the mine, shadow, mine background, and shadow background regions.

D. Classifier Construction

This section presents construction of a binary classifier for discrimination of mine-like and nonmine-like objects by using the four features (stochastic matrices) corresponding to different regions of the geometric model, as described in the previous section. In addition, two stochastic matrices $\bar{\Pi}_{Mine}$ and $\bar{\Pi}_{Shadow}$ are defined as the expected (or averaged) stochastic signatures of mines and shadows, respectively. These are estimated from known mine and shadow regions in the training set. The expected stochastic behavior of the background regions are not used as they exhibit very large variations. Four scalar measures,

η^m , η^b , η^{mb} , and η^{sb} , that are dependent on the pixel location (i, j) , are derived as follows:

$$\eta^m(i, j) \triangleq d(\Pi^m(i, j), \bar{\Pi}_{Mine}) \quad (5)$$

$$\eta^s(i, j) \triangleq d(\Pi^s(i, j), \bar{\Pi}_{Shadow}) \quad (6)$$

$$\eta^{mb}(i, j) \triangleq d(\Pi^{mb}(i, j), \bar{\Pi}_{Mine}) \quad (7)$$

$$\eta^{sb}(i, j) \triangleq d(\Pi^{sb}(i, j), \bar{\Pi}_{Shadow}) \quad (8)$$

where $d(\bullet, \bullet)$ is a valid distance measure between two stochastic matrices. Possible candidates for $d(\bullet, \bullet)$ include l_1 , l_2 distance, induced operator norm (l_1 or l_2) or the (possibly symmetrized) Kullback–Leibler divergence.

The scalar measure $\eta^m(i, j)$ corresponds to the deviation of the stochastic matrix $\Pi^m(i, j)$ from the expected mine signature $\bar{\Pi}_{Mine}$. This deviation is expected to be small for an actual mine; however, it could be large for nonmines. Similarly, $\eta^s(i, j)$ denotes the deviation of the $\Pi^s(i, j)$ from the expected shadow signature $\bar{\Pi}_{Shadow}$, which is obtained from the region A^s . As in the previous case, smaller values of $\eta^s(i, j)$ indicate the presence of shadows. Furthermore, $\eta^{mb}(i, j)$ is defined as the distance of the mine background $\Pi^{mb}(i, j)$ from $\bar{\Pi}_{Mine}$; and $\eta^{sb}(i, j)$ is defined as the distance of the shadow background $\Pi^{sb}(i, j)$ from $\bar{\Pi}_{Shadow}$. The primary importance of using background as a feature in identifying mine-like objects is to distinguish mines from other large-size objects (e.g., ship wreckage, rocky bottoms), and artifacts (e.g., those generated from vehicle roll, vehicle turns and surface returns). Mine-like-objects usually exhibit greater dissimilarities between the mine (or shadow) regions and the mine background (or shadow background) region.

A threshold-based sequential testing scheme has been formulated to classify the observed patterns of the sonar images into mine-like and nonmine-like categories. Given appropriately chosen scalar thresholds λ_1 , λ_2 , λ_3 , and λ_4 , the following sequential tests must be satisfied for a pixel (i, j) to be classified as part of a mine.

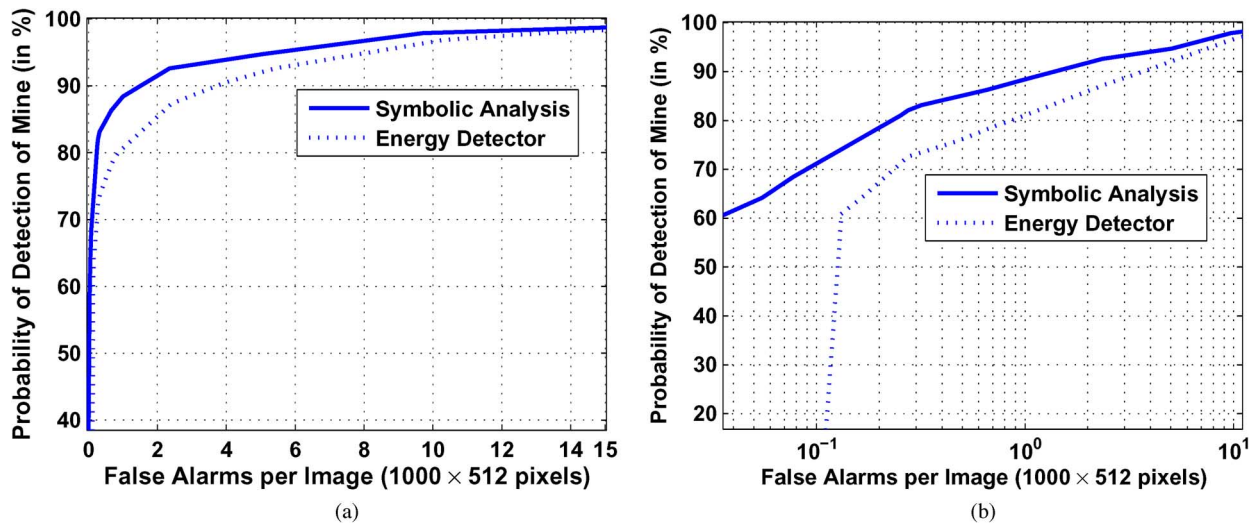


Fig. 8. Comparison of receiver operating characteristics (ROC) for mine detection using symbolic analysis and energy detector. (a) Both FAR and P_D in the linear scale; (b) FAR in the log scale and P_D in the linear scale.

- 1) $\eta^m(i, j) \leq \lambda_1$: This confirms that the stochastic matrix $\Pi^m(i, j)$ closely matches that of a mine.
- 2) $\eta^s(i, j) \leq \lambda_2$: This confirms that the stochastic matrix $\Pi^s(i, j)$ closely resembles that of a shadow.
- 3) $\eta^{mb}(i, j) \geq \lambda_3$ or $\eta^{sb}(i, j) \geq \lambda_4$: This confirms that the backgrounds adjacent to the mine or the shadow are sufficiently distinct from the mine or the shadow themselves.

A binary decision is made to determine whether the pixel location belongs to a potential mine-like object. A number (1 or 0) is assigned to each pixel of the image based on the classification as a mine-like or nonmine-like object, respectively. Subsequently, a smoothing algorithm is applied to remove the outliers such that isolated 1's (that are possibly false alarms within a region where no other point in the neighborhood is detected as a mine) are removed. If several 1's occur close to each other, then the whole region is clustered together and counted as a single detection. A disk like structure is used for morphological erosion operation [32] via the MATLAB command "imerode." In summary, the algorithm provides a binary decision map for every input sonar image.

The above four scalar thresholds, i.e., λ_1 , λ_2 , λ_3 , and λ_4 , are chosen based on the receiver operating characteristics (ROC) [33] in Fig. 8 to yield different choices of the probability of detection (P_D) and the false alarm rate (FAR), which are defined in terms of percentage of correctly detected mines and the number of mines falsely detected per image, respectively. (Note: the size of an image is 512 × 1000 pixels.)

Remark 2: The occurrence of mines in sonar images is sparse, resulting in orders of magnitude higher amount of data for nonmines (e.g., background) as compared to mines. Therefore, conventional classification techniques such as support vector machines (SVM) [23] have not been used because of the high degree of imbalance in the training data set.

E. ROC

To construct the receiver operating characteristics (ROC), a training data set consisting of 91 images (containing a total of 100 mines) is considered. These images consist of various

textured backgrounds, with different types of seabed objects, sand-ripples, and rock outcrops. The image quality is altered because of distortion induced by vehicle motion and artifacts caused by beam pattern imperfections. The images (512 × 1000 pixels), in the range of 0 to 255 on the gray scale, are partitioned to create a symbolized image. A sliding window method is used to implement the geometric model in Section II and the classifier described above. For each pixel in the image, a model is constructed, as shown in Fig. 7, such that the pixel under consideration is at the center of the mine region of the geometric model. The PFSA are generated from the four regions (mine, shadow, mine background and shadow background) based on interactions between symbols at neighboring pixel locations ($\kappa = 1$) and four features (stochastic matrices) Π^m , Π^s , Π^{mb} and Π^{sb} are obtained.

Remark 3: The images analyzed in this paper are obtained from low-resolution long-range sidescan sonar data. In these images, the mines appear as extremely small objects, often as wide as only 3–5 pixels. As a result, obtaining sufficient data to populate the state transition matrix of the PFSA depends upon the number of states. In order to keep the number of states under check, the neighborhood system with $\kappa = 1$ has been chosen because $\kappa > 1$ would significantly increase the number of PFSA states, for which sufficient data may not be available (e.g., in the exploration phase of mine hunting). However, $\kappa > 1$ could be used for high-resolution sonar images that capture mine-like-objects in greater detail when more scan data are available (e.g., in the exploitation phase of mine hunting).

For the purpose of feature extraction, the distance measure between the stochastic matrices is chosen as the l_1 norm. Each of the four thresholds (i.e., λ_1 , λ_2 , λ_3 , and λ_4), defined in Section III-D, are varied from 0 to the maximum possible distance between two stochastic matrices (which is 6 for the choice of l_1 norm) in steps of 0.5. For a certain combination of parameters, the classification algorithm is executed over the entire training data set and the number of false alarms and the number of correct detections are counted to obtain the false alarm rate (FAR) per image and the probability of detection

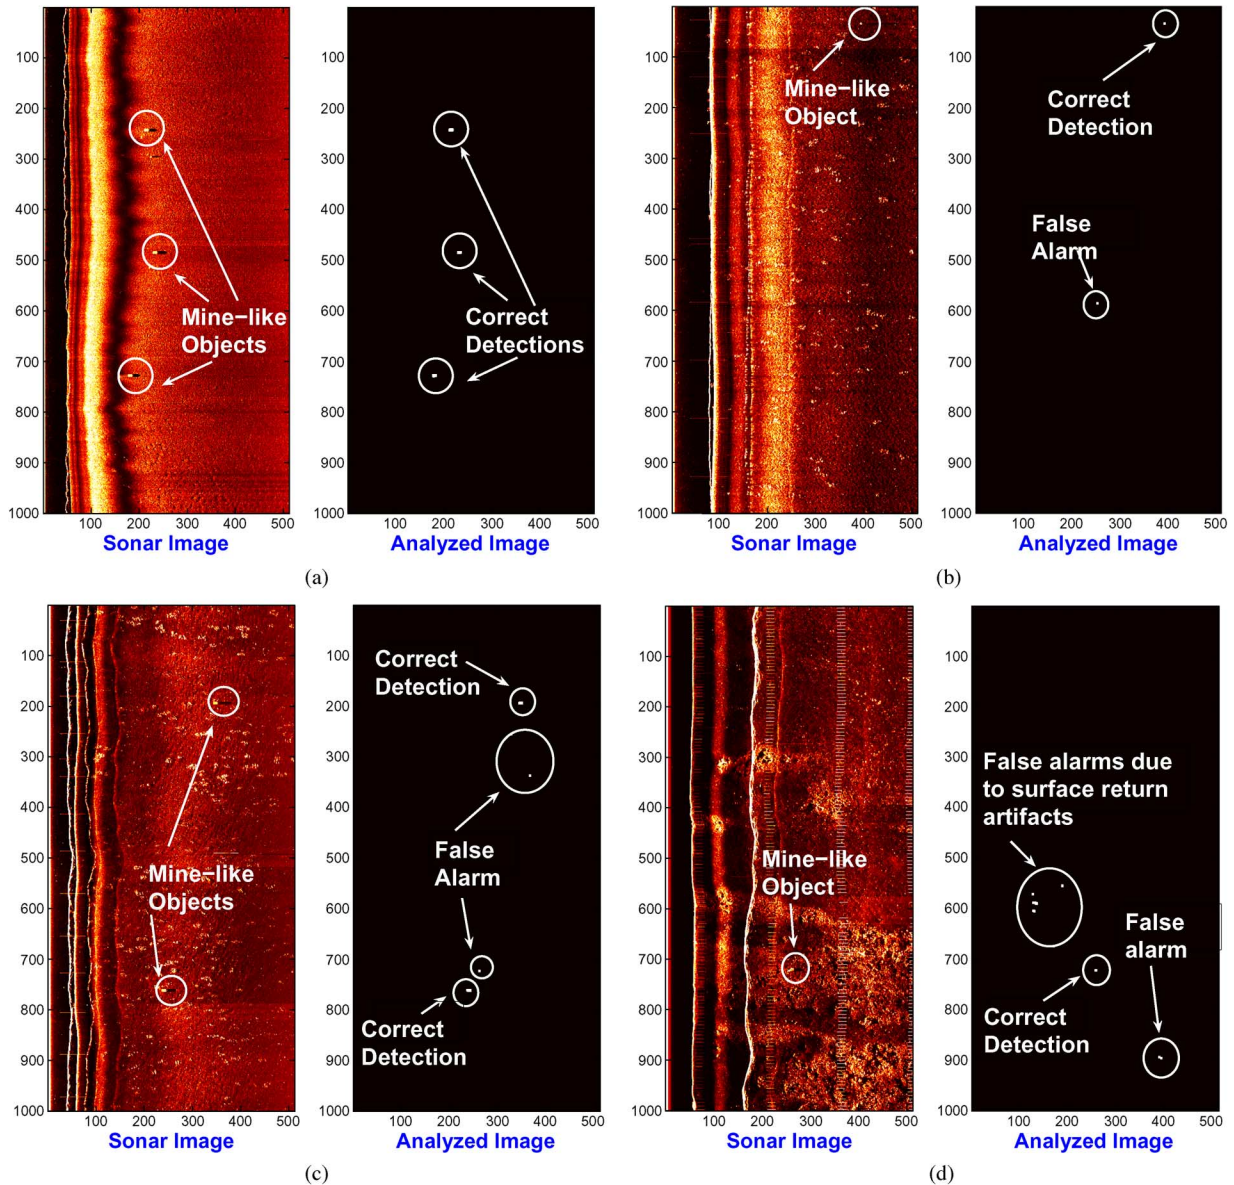


Fig. 9. Mine detection from representative images in the test data set with different levels of noise (e.g., reverberation, clutter, and artifacts). (a) Symbolic analysis of a relatively clean sonar image; (b) symbolic analysis of a sonar image with moderate noise; (c) symbolic analysis of a sonar image with high noise; (d) symbolic analysis of a sonar image with very high noise.

(P_D), respectively. A nondominated genetic algorithm [34], [35] is used to progressively select the combinations of the four thresholds to generate a population of feasible points. The ROC plot is constructed by joining the outermost points on the plot of P_D versus FAR as shown in Fig. 8. Therefore, every point on the ROC maps to a 4-tuple of threshold values.

In practice, the four thresholds λ_1 , λ_2 , λ_3 , and λ_4 are chosen corresponding to a desired operating point on the ROC. For subsequent analysis, the thresholds were chosen as $\lambda_1 = 2.5$, $\lambda_2 = 3.5$, $\lambda_3 = 4.0$ and $\lambda_4 = 5.5$, which correspond to a probability of detection of 90% and 1.3 false alarms per image, where each image covers an area of approximately $30 \text{ m} \times 30 \text{ m}$.

For the purpose of comparison, the same sets of images have been analyzed based on the data generated by an energy detector. This detector implements individual thresholds to separate both the highlight and the shadow from the background

reverberation and clutter. The energy detector is used in conjunction with the same geometric model, described in Section II, to depict the utility of symbolic pattern analysis. A comparison of the ROCs as obtained by symbolic analysis and energy detector are depicted in Fig. 8. In Fig. 8(a), both FAR and P_D are in the linear scale; and Fig. 8(b) shows the same ROCs with FAR in the log scale and P_D in the linear scale. These ROC plots depict the advantage gained in successful mine detection with relatively low false alarm rates.

IV. RESULTS AND DISCUSSION

This section presents the results generated upon execution of the symbolic pattern analysis algorithm on 60 images from the test data set that is different from the training data set used to generate the partition, the matrices $\bar{\Pi}_{Mine}$ and $\bar{\Pi}_{Shadow}$, and the ROC. As an example, four test images are shown in Fig. 9 with different levels of induced noise contamination, where the

TABLE I
RESULTS OF SYMBOLIC PATTERN ANALYSIS ON A TEST DATA SET

Description	Value
Number of Test Images Analyzed	60
Probability of Detection (P_D)	(53/58) > 90%
False alarm rate (FAR) per image	~1.5

results of detection are shown in the right hand side of each plot from (a) to (d). The analyzed images are binary and consist of only two symbols 0 (i.e., absence of a mine) and 1 (i.e., presence of at least one mine).

In the testing of the algorithm, the mine location reported is expected to differ to some extent from what is calculated in the ground truth. This problem is addressed by defining a “tolerance radius” [10] that allows a detected object to have its centroid within the neighborhood of a ground-truth object with this tolerance radius. Mine detections that are outside these neighborhoods are counted as false alarms; and the mines that are never detected are counted as missed detections. The tolerance radius is a characteristic of the objects being sought and the sonar being used. The tolerance radius is chosen here as the size of an average mine in the ground-truth file.

Four representative images are shown in Fig. 9, each showing different levels of background noise consisting of seabed reverberation and clutter. A representative set of threshold values (see Section III-E) is chosen to yield a high detection probability with an acceptable false alarm rate based on the premise that a missed detection costs much more than a false alarm. Tests show that the algorithm is capable of detecting mines in a high concentration of seabed clutter, including mines buried under vehicular artifacts. The algorithm has been executed on the entire set of test data with the same values of representative thresholds; Table I lists false alarms and successful detections with 53 successful detections out of a total of 58 mines (i.e., $P_D > 90\%$), while the false alarm rate (FAR) is approximately 1.5 per image (1000×512 pixels). With alternative choices of operating points on the ROC curves, the mine detection probability can be increased at a modest expense of increased FAR .

Remark 4: The algorithm presented in this paper has been applied to the sonar images without any preprocessing. Apparently, the performance (e.g., reduced FAR) could be further improved if adaptive clutter suppression filters, such as those described in [36], are used in the preprocessing stage.

Remark 5: In practice, the proposed algorithm could be used in conjunction with other state-of-the art techniques for the purpose of mine detection. Information fusion algorithms that combine detections from various classifiers have been shown to yield a better ROC [37], [38]. The decisions derived from the proposed algorithm, could benefit the overall performance if they are made in conjunction with information fusion, especially because this algorithm is based on significantly different principles from other reported algorithms such as those used in [39].

V. SUMMARY, CONCLUSION, AND FUTURE WORK

This paper presents a symbolic pattern analysis method for detection of underwater mines using sonar data in the form of

images. Underlying algorithms of pattern analysis are formulated based upon the concepts derived from *symbolic dynamics* and *automata theory*. Specifically, the paper presents a methodology to construct probabilistic finite state automata (PFSA) from two-dimensional (i.e., image) data sets that have been symbolized using a physics-based partitioning scheme. These PFSA represent the underlying features of data such that their probabilistic structure (i.e., the state transition matrix) evolve at different spatial locations in the image.

A key aspect of this paper is construction of a geometric model for mine detection in sonar images. Different regions in the model are carefully chosen to match the sizes of the objects and other relevant features (i.e., shadow and clutter) that are to be detected. The observed patterns corresponding to different regions of the geometric model are then fused in a classification scheme to make binary decisions as mine-like and nonmine like categories, especially in the exploration phase of mine hunting operation when the available data are limited and low resolution.

The algorithm has been tested on different images of 512×1000 pixels (900 m^2). The probability of correct detection of mines is found to be ~90% with an average of ~1.5 false alarms in each image. These results show that the algorithm is capable of detecting underwater mines with a low probability of false alarms in different images that have varying degrees of noise and seabed clutter.

The major advantages of the proposed pattern analysis algorithm for underwater mine detection are delineated below.

- 1) Performance of the statistical symbolic pattern analysis algorithm is robust with respect to both noise and echo artifacts in sonar images because of coarse graining of the sonar data. Furthermore, this algorithm is robust to variations in the mine characteristics such as size and length of the shadow provided that they are within specified bounds. Performance of the algorithm has also been tested to be robust with respect to the locations of segment boundaries of the partition. The important aspect of the physics-based partition is that it corresponds to the three characteristic features, namely, mine, shadow, and the background.
- 2) The algorithm is computationally efficient in terms of execution time and memory requirements as a consequence of small alphabet size and a small number of states in probabilistic automata. As such, the algorithm can be programmed and powered on small microprocessors and implemented on-board for real-time application on an unmanned underwater vehicle (UUV).
- 3) In contrast to traditional Bayesian methods such as the likelihood-ratio-test, the symbolic pattern analysis method does not require *a priori* knowledge of probability distributions for characterizing mines and nonmine-like objects. Specifically, because of the design of geometric models, the proposed algorithm is applicable even if the unknown probability distributions are multimodal.

Further theoretical and experimental research is necessary before the proposed algorithm could be considered for implementation in the ocean environment. While there are many such issues in this regard, the following topics are under active research:

- application of cross validation techniques to verify the robustness of the classifier;
- testing under real-life scenarios that include varying bathymetric properties, various ocean depths and different sea states;
- testing of the algorithm performance for simultaneous enhancement of successful detection and reduction of false alarms through additional measurements such as multiple scanning from different angles;
- improvement of image segmentation into three symbols $\{a, b, c\}$ through usage of noise distribution models as described in [40];
- enhancement of algorithm performance through usage of flexible and adaptive geometric models (e.g., varying shadow lengths) to account for changes in the height above bottom for the detector in addition to implementing a range dependant geometric model;
- extension of the present algorithm with $\kappa = 1$ for limited data (e.g., in the exploration phase of a mine hunting operation) to that with $\kappa > 1$ for high-resolution sonar images to capture mine-like-objects in greater detail when more scan data are available (e.g., in the exploitation phase);
- comparison of the proposed algorithm with other detection methods such as those described in [41].

ACKNOWLEDGMENT

The authors would like to thank the Panama City Division of the Naval Surface Warfare Center, Panama City, FL, for providing the data used in this publication. The authors would also like to thank the assistance of their colleague, Dr. C. Rao, in this work.

REFERENCES

- [1] D. Li, K. Wong, Y. Hu, and A. Sayeed, "Detection, classification, and tracking of targets," *IEEE Signal Proc. Mag.*, pp. 17–29, Mar. 2002.
- [2] R. J. Urick, *Principles of Underwater Sound*, 3rd ed. New York: McGraw-Hill, 1983.
- [3] F. Crosby and J. Cobb, "Sonar processing for short range, very-high resolution autonomous underwater vehicle sensors," in *Proc. MTS/IEEE OCEANS*, 2005, vol. 1, pp. 398–402.
- [4] R. E. Hansen, T. O. Saebo, K. Gade, and S. Chapman, "Signal processing for AUV based interferometric synthetic aperture sonar," in *Proc. MTS/IEEE OCEANS*, San Diego, , September 2003.
- [5] R. E. Hansen, J. Groen, and H. J. Callow, "Image enhancement in synthetic aperture sonar," in *Proc. Int. Conf. Detection Classif. Underwater Target.*, 2007, vol. 29, no. 6, pp. 69–76.
- [6] B.-G. Kim, D.-J. Kim, and D.-J. Park, "Novel precision target detection with adaptive thresholding for dynamic image segmentation," *Mach. Vis. Appl.*, vol. 12, pp. 259–270, 2001.
- [7] S. Reed, Y. Petillot, and J. Bell, "Automated approach to classification of mine-like objects in sidescan sonar using highlights and shadow information," *IEE Proc.-Radar Sonar Navig.*, vol. 151, no. 1, pp. 90–105, Feb. 2004.
- [8] E. Dura, Y. Zhang, X. Liao, G. Dobeck, and L. Carin, "Active learning for detection of mine-like objects in side-scan sonar imagery," *IEEE J. Ocean. Eng.*, vol. 30, no. 2, pp. 360–371, Apr. 2003.
- [9] A. H. Gunatilaka and B. A. Baertlein, "Feature-level and decision-level fusion of noncoincidentally sampled sensors for land mine detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 6, pp. 577–589, Jun. 2001.
- [10] B. Calder, L. Linnett, and D. Carmichael, "Bayesian approach to object detection in sidescan sonar," *IEE Proc. Vis. Image Signal Process.*, vol. 145, no. 3, pp. 221–228, 1998.
- [11] F. Maussang, J. Chanussot, A. Hétet, and M. Amate, "Higher order statistics for the detection of small objects in a noisy background, application on sonar imaging," *EURASIP J. Adv. Signal Process.*, vol. 2007, 10.1155/2007/47039.
- [12] P. Bharadwaj and L. Carin, "Infrared-image classification using hidden markov trees," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 10, pp. 1394–1398, Oct. 2002.
- [13] A. Quattoni, S. Wang, L. Morency, M. Collins, and T. Darrell, "Hidden conditional random fields," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 10, pp. 1848–1853, Oct. 2007.
- [14] X.-F. Ye, Z.-H. Zhang, P. X. Liu, and H.-L. Guan, "Sonar image segmentation based on GMRF and level-set models," *Ocean Eng.*, vol. 37, no. 10, pp. 891–901, 2010.
- [15] F. Maussang, J. Chanussot, A. Hétet, and M. Amate, "Mean-standard deviation representation of sonar images for echo detection: Application to sas images," *IEEE J. Ocean. Eng.*, vol. 32, no. 4, pp. 956–970, Oct. 2007.
- [16] F. Maussang, M. Rombaut, J. Chanussot, A. Hétet, and M. Amate, "Fusion of local statistical parameters for buried underwater mine detection in sonar imaging," *EURASIP J. Adv. Signal Process.*, vol. 2008, 10.1155/2008/876092.
- [17] H. M. Dee and S. A. Velastin, "How close are we to solving the problem of automated visual surveillance?," *Mach. Vis. Appl.*, vol. 19, no. 5–6, pp. 329–343, 2008.
- [18] A. Ray, "Symbolic dynamic analysis of complex systems for anomaly detection," *Signal Process.*, vol. 84, no. 7, pp. 1115–1130, 2004.
- [19] V. Rajagopalan and A. Ray, "Symbolic time series analysis via wavelet-based partitioning," *Signal Process.*, vol. 86, no. 11, pp. 3309–3320, 2006.
- [20] S. Gupta and A. Ray, "Statistical mechanics of complex systems for pattern identification," *J. Statist. Phys.*, vol. 134, no. 2, pp. 337–364, 2009.
- [21] D. Lind and M. Marcus, *An Introduction to Symbolic Dynamics and Coding*. Cambridge, U.K.: Cambridge Univ. Press, 1995.
- [22] R. Pathria, *Statistical Mechanics*, 2nd ed. Oxford, U.K.: Butterworth-Heinemann, 1996.
- [23] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. New York: Wiley Interscience, 2001.
- [24] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 1st ed. New York: Wiley Interscience, 1991.
- [25] H. E. Hopcroft, R. Motwani, and J. D. Ullman, *Introduction to Automata Theory, Languages, and Computation*, 2nd ed. Boston, MA: Addison Wesley, 2001.
- [26] S. Gupta, A. Ray, and E. Keller, "Symbolic time series analysis of ultrasonic data for early detection of fatigue damage," *Mech. Syst. Signal Process.*, vol. 21, no. 2, pp. 866–884, 2007.
- [27] S. Gupta and A. Ray, "Real-time fatigue life estimation in mechanical systems," *Measurement Sci. Technol.*, vol. 18, no. 7, pp. 1947–1957, 2007.
- [28] G. Mallapragada, I. Chattopadhyay, and A. Ray, "Automated behavior recognition in mobile robots using symbolic dynamic filtering," *J. Syst. Contr. Eng.*, 2008.
- [29] C. Rao, A. Ray, S. Sarkar, and M. Yasar, "Review and comparative evaluation of symbolic dynamic filtering for detection of anomaly patterns," *Signal, Image, Video Process.*, vol. 3, no. 2, pp. 101–114, 2009.
- [30] M. Buhl and M. Kennel, "Statistically relaxing to generating partitions for observed time-series data," *Phys. Rev. E*, vol. 71, no. 4, p. 046213, 2005.
- [31] G. R. Grimmett, "A theorem about random fields," *Bull. London Math. Soc.* vol. 5, no. 1, pp. 81–84, 1973 [Online]. Available: <http://blms.oxfordjournals.org>
- [32] G. Ginolhac, J. Chanussot, and C. Hory, "Morphological and statistical approaches to improve detection in the presence of reverberation," *IEEE J. Ocean. Eng.*, vol. 30, no. 4, pp. 881–899, 2005.
- [33] V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. Berlin, Germany: Springer-Verlag, 1988.
- [34] I. Das and J. Dennis, "Normal-boundary intersection: A new method for generating the pareto surface in nonlinear multicriteria optimization problems," *SIAM J. Optim.*, vol. 84, no. 7, pp. 631–657, 1998.
- [35] T. Wettergren, The Genetic-Based Normal Boundary Intersection (GANBI) Method: An Efficient Approach to Pareto Multiobjective Optimization for Engineering Design Naval Undersea Warfare Center Division, Newport, RI, NUWC Technical Report 11, 741, May 15, 2006.

- [36] T. Aridgides, M. F. Fernandez, and G. J. Dobeck, Adaptive Three-Dimensional Range-Crossrange-Frequency Filter Processing String for Sea Mine Classification in Side Scan Sonar Imagery vol. 3079, no. 1, SPIE, 1997, pp. 111–122.
- [37] G. Dobeck, “Algorithm fusion for automated sea mine detection and classification,” in *Proc. MTS/IEEE Conf. Exhib. OCEANS*, 2001.
- [38] C. M. Ciany, W. Zurawski, and G. J. Dobeck, Application of Fusion Algorithms for Computer-Aided Detection and Classification of Bottom Mines to Shallow Water Test Data From the Battle Space Preparation Autonomous Underwater Vehicle (BPAUV) vol. 5089, no. 1, SPIE, 2003, pp. 58–64.
- [39] G. Dobeck, “A probabilistic model for score-based algorithm fusion,” in *Proc. MTS/IEEE OCEANS*, 2005.
- [40] M. Mignotte, C. Collet, P. Prez, and P. Bouthemy, “Three-class markovian segmentation of high-resolution sonar images,” *Comput. Vis. Image Understand.*, vol. 76, no. 3, pp. 191–204, 1999.
- [41] T. Aridgides, M. Fernandez, and G. Dobeck, “Fusion of sea mine detection and classification processing strings for sonar imagery,” in *Proc. SPIE 4038*, 2000, pp. 391–401.



Kushal Mukherjee (S'09) received the B.Sc. degree in mechanical engineering from the Indian Institute of Technology (IIT), Roorkee, India, in 2006. He received dual M.Sc. degrees in mechanical engineering and electrical engineering from the Pennsylvania State University, University Park, in 2009, where he is currently working towards the Ph.D. degree in mechanical engineering.

His research interests include symbolic dynamics, sensor networks, and multiagent systems. He is a student member of the American Society of Mechanical

Engineers (ASME).



Shalabh Gupta (S'05–M'07) received dual M.Sc. degrees in mechanical engineering and electrical engineering in 2004 and 2005, as well as the Ph.D. degree in mechanical engineering from the Pennsylvania State University, University Park, in 2006.

He is currently a Research Associate in the Department of Mechanical Engineering at the Pennsylvania State University. His research interests include the science of autonomy, swarm robotics, intelligent systems, machine learning, network science, and fault detection and isolation in complex systems.

His research efforts have been instrumental in opening new fields of data understanding and pattern discovery via interfacing multidisciplinary concepts derived from statistical mechanics, symbolic dynamics, and languages and automata theory.

Dr. Gupta is a member of the American Society of Mechanical Engineers (ASME) and the Institute of Electrical and Electronics Engineers (IEEE).



Asok Ray (SM'83–F'02) received graduate degrees in electrical engineering, mathematics, and computer science in 1970, 1978, and 1972, respectively. He also received the Ph.D. degree in mechanical engineering from Northeastern University, Boston, MA, in 1976.

He joined the Pennsylvania State University, University Park, in July 1985, and is currently a Distinguished Professor of Mechanical Engineering and a Graduate Faculty of Electrical Engineering. Prior to joining Penn State, he also held research and academic positions at Massachusetts Institute of Technology and Carnegie-Mellon University, as well as research and management positions at GTE Strategic Systems Division, Charles Stark Draper Laboratory, and MITRE Corporation. His research interests include control and optimization of continuously varying and discrete-event dynamical systems, intelligent instrumentation for real-time distributed systems, and modeling and analysis of complex dynamical systems from thermodynamic perspectives in both deterministic and stochastic settings, as applied to robotics, undersea autonomous vehicles, aeronautics and astronautics, and fossil-fueled and nuclear power plants. He has authored or coauthored over 450 research publications including about 250 scholarly articles in refereed journals and research monographs.

Dr. Ray had been a Senior Research Fellow at NASA Glenn Research Center under a National Academy of Sciences award. He is a Fellow of ASME and a Fellow of World Innovative Foundation (WIF). Further details of Dr. Ray's credentials are available in the web address <http://www.mne.psu.edu/Ray/>



Shashi Phoha (M'91–SM'91) received dual graduate degrees in operation research and mathematics in 1972 and 1970, respectively, and the Ph.D. degree in mathematical statistics in 1976.

She is currently a Professor of Electrical Engineering and the Director of the Information Science and Technology Division of the Applied Research Laboratory at Penn State University, University Park. She has led multiorganizational advanced research programs and laboratories in major U.S. industrial and academic institutions. She was the

Director of the Information Technology Laboratory of the National Institute of Standards and Technology. Her research focuses on information sciences, which bring together ideas from logic, operations research, formal languages, automata theory, dynamic systems, and control theory in diverse applications.

Dr. Phoha was awarded the IEEE Technical Achievement Award in 2005 by the IEEE Computer Society. Drawing on her multidisciplinary research innovations, she pioneered the use of formal methods for the scientific analysis of distributed information for decision support, multistage coordination and intelligent control. She has been an Associate Editor of the IEEE International Conference on Systems, Man, and Cybernetics, and is an Editor of the *International Journal of Distributed Sensor Networks*.