Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE DIRECT®

Signal Processing 84 (2004) 1225 – 1229

SIGNAL PROCESSING

www.elsevier.com/locate/sigpro

Anomaly detection based on an iterative local statistics approach☆

Arnon Goldman, Israel Cohen*

Department of Electrical Engineering, Technion, Israel Institute of Technology, Technion City, Haifa 32000, Israel

Received 26 September 2003; received in revised form 23 March 2004

Abstract

One of the most challenging problems in automatic target detection is associated with the large variability of background clutter and object appearance. In this paper, we propose an anomaly detection approach which does not rely on an exhaustive statistical model of the targets, but rather on the local statistics of the data and possibly on some a priori information regarding the sizes and shapes of targets. Iterative procedures of feature extraction and anomaly detection are carried out, gradually reducing the false alarm rate while maintaining a high probability of detection. The background is characterized in a feature space of principal components, and a single hypothesis scheme is used for the detection of anomalous pixels. Morphological operators are subsequently employed for extracting the sizes and shapes of anomalous clusters in the image domain, and identifying potential targets. The robustness of the proposed approach is demonstrated with application to sea-mine detection in sonar imagery.

Keywords: Anomaly detection; Object recognition; Pattern recognition; Image segmentation; Image texture analysis

1. Introduction

Target detection in radar and sonar imagery is a challenging problem due to the large variability in background clutter and in object appearance. The detection of sea-mines, for example, involves addressing the varying shape of the ocean surface and its vegetation [2]. Land-mine detection, using ground penetrating radar (GPR) encounters the problem of the extreme clutter environment within the first 5 cm of the soil surface. Almost any object under the surface of the ground yields a return signal, which may be confused with a lethal land-mine [6]. In most cases, lethal targets must be detected with nearly 100% reliability. False detections may not be disastrous but might slow down the dimining process.

The majority of work in the area of target detection has focused on detection methods, which involve statistical characterization of both targets and background [9]. Classification methods which are based on the Bayesian approach, require knowledge of a priori statistics (i.e., conditional density functions) that may be a source of loss of robustness. Matched filters, for example, require a priori knowledge of a typical signature of the target [2]. In a realistic situation, however, there is a wide variety of

☆ This research was supported by the fund for the promotion of research at the Technion.

* Corresponding author.
E-mail address: icohen@ee.technion.ac.il (I. Cohen).
URL: http://www-sipl.technion.ac.il/∼cisrael

0165-1684/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.sigpro.2004.04.004

potential targets which do not conform to a uniform model.

In this paper, we propose an anomaly detection approach, which does not rely on an exhaustive statistical model of the targets, but rather on the local statistics of the data and possibly on some a priori information regarding the sizes and shapes of targets. An iterative procedure of feature extraction, based on local statistics and principal components analysis (PCA) is performed. The background is statistically characterized in a feature space of principal components. A single hypothesis scheme is used for the detection of anomalous pixels in a given region of interest (ROI). Subsequently, morphological operators are employed for extracting the sizes and shapes of anomalous clusters in the image domain, and identifying potential targets. We may compromise on the false alarm rate in order to achieve a high probability of detection, since each iteration gradually reduces the false alarm rate while maintaining the high probability of detection.

The paper is organized as follows. Section 2 presents a formulation of the problem, Section 3 provides a mathematical model, Section 4 describes the proposed anomaly detection algorithm, and Section 5 demonstrates the application of the algorithm to sea-mine detection in sonar imagery.

2. Problem formulation

Let Ω be the support of a gray-scale image. The image may contain different textures as background (i.e. grass, trees, soil, water) and some targets which are anomalous with regards to the background. We would like to find a disjoint partition of Ω, such that $\Omega = \mathscr{B} \cup \mathscr{A}$, where $\mathscr{B}$ contains a few uniform subsets of background pixel clusters: $\mathscr{B} = \bigcup_{k=1}^{K} \mathscr{B}_k$, and $\mathscr{A}$ contains anomalous pixels. The subsets of the background pixels represent different textures of the background. For example, in a typical terrain scene, $\mathscr{B}_1$ may contain pixels which constitute grass, $\mathscr{B}_2$ may represent trees, and so on.

The proposed algorithm is based on an iterative two-category classification procedure. A set of pixels, $\mathscr{A}_{k-1}$, which have been identified as anomalous in the $(k-1)$th iteration, is partitioned into background and anomalous subsets, $\mathscr{B}_k$ and $\mathscr{A}_k$, respectively

$$\mathscr{A}_{k-1} = \mathscr{B}_k \cup \mathscr{A}_k, \tag{1}$$

where $k \geqslant 1$ and $\mathscr{A}_0$ is initialized to the ROI in the image (possibly $\mathscr{A}_0 = \Omega$).

The fundamental problem of the iterative anomaly detection procedure is to partition the set of pixels, $\mathscr{A}_{k-1}$, into two subsets: one, which is relatively uniform, large and thus considered as the background subset ($\mathscr{B}_k$), and a second subset which contains anomalous pixels ($\mathscr{A}_k$). Another problem is to determine a stopping rule for the iterative procedure. The procedure should be iterated until the number of targets pertaining to $\mathscr{A}_k$ does not exceed a given number. The maximal number of detectable targets in a given ROI is closely related to the permissible false alarm rate. A stopping rule based on such a criterion should incorporate into the procedure a low complexity routine for coarse target detection.

3. Mathematical model

Let $\lambda \in \Omega$ be the coordinates of a pixel in the image. Let q_λ be a feature vector related to the pixel λ. Considering a two-category classification problem, we define two possible hypotheses:

$$H_0 : \lambda \in \mathscr{B} \quad \text{(background pixel)}$$

$$H_1 : \lambda \in \mathscr{A} \quad \text{(anomalous pixel)}.$$

A two-category classification problem is often worked out using the Bayes decision rule for minimum cost [1]. Let $C(H_1|H_0)$ and $C(H_0|H_1)$ denote, respectively, the costs of false detection and miss detection. Then the optimal minimum cost decision rule is given by [5]

$$f(q_\lambda|H_0)P(H_0)C(H_1|H_0) \underset{H_1}{\overset{H_0}{\gtrless}} f(q_\lambda|H_1)P(H_1)C(H_0|H_1), \tag{2}$$

where $f(q_\lambda|H_i)$ is the conditional pdf of q_λ given H_i and $P(H_i)$ is the a priori probability of H_i.

In practice, the background can be well characterized (an empirical $f(q_\lambda|H_0)$ can be generated) while the anomalies are of a wide variety and rare (a reliable estimate of $f(q_\lambda|H_1)$ is unavailable). Therefore, the Bayes decision rule for minimum cost is inapplicable. A suitable alternative to such problems is based on a single hypothesis scheme [5]. Let $\mu = E[q_\lambda|H_0]$ denote the expected feature vector and $\Sigma = E[(q_\lambda - \mu)(q_\lambda - \mu)^{\mathrm{T}}|H_0]$ the covariance matrix under H_0 hypothesis. Let the normalized distance of q_λ from its expected vector, μ, be defined by

$$d(q_\lambda) = (q_\lambda - \mu)^{\mathrm{T}}\Sigma^{-1}(q_\lambda - \mu). \tag{3}$$

Then the decision rule is given by

$$d(q_\lambda) \underset{H_1}{\overset{H_0}{\lessgtr}} D, \tag{4}$$

where D is the threshold to determine whether a given pixel is anomalous or not. This decision rule is based on the statistics of the background only. No knowledge about the anomalies statistics is taken into consideration. The threshold, D, can be determined according to a specified confidence level, η, which is the probability of correctly deciding on H_0 given H_0 is true. The threshold, D, and the confidence level, η, are related by

$$\eta \equiv \Pr(H_0|H_0) = \Pr(d(q_\lambda) \leqslant D|H_0). \tag{5}$$

In case the feature vector, q_λ, is a Gaussian random vector of dimension n, the pdf of $d^2(q_\lambda)$ under the H_0 hypothesis, denoted by $p_{d^2}(\zeta)$, is the gamma density function with parameters $\beta = n/2 - 1$ and $\alpha = 1/2$ [5]. Accordingly, the relation between η and D can be written as

$$\eta = \int_0^{D^2} p_{d^2}(\zeta)\,\mathrm{d}\zeta$$
$$= \int_0^{D^2} \frac{1}{2^{n/2}\Gamma(n/2)}\,\zeta^{(n-2)/2}\mathrm{e}^{-\zeta/2}\,\mathrm{d}\zeta. \tag{6}$$

The aforementioned decision rule is based on a one-dimensional measure, $d(q_\lambda)$, whereas the feature vector, q_λ is multi-dimensional (size n). It is quite clear that reducing the dimension of the feature space to a one-dimensional distance space eliminates valuable information. Accordingly the false detection rate is likely to increase. However high false detection rate is acceptable, since we consider an iterative procedure, and each iteration gradually reduces the false alarm rate while maintaining a high probability of detection.

4. Anomaly detection algorithm

A block diagram of the proposed algorithm is presented in Fig. 1. The algorithm consists of two iterative procedures. The basic procedure is an iterative disjoint partition of $\mathscr{A}_{k-1}$ into $\mathscr{A}_k$ and $\mathscr{B}_k$. After the partitioning, the decision whether to further partition $\mathscr{A}_k$ depends on the number of potential targets found in $\mathscr{A}_k$. This procedure iterates, reducing $\mathscr{A}_k$'s population ($\mathscr{A}_k \subset \mathscr{A}_{k-1}$) until the number of potential targets, found in $\mathscr{A}_k$, is smaller than or equal to a given number, N. The second procedure is implicit in

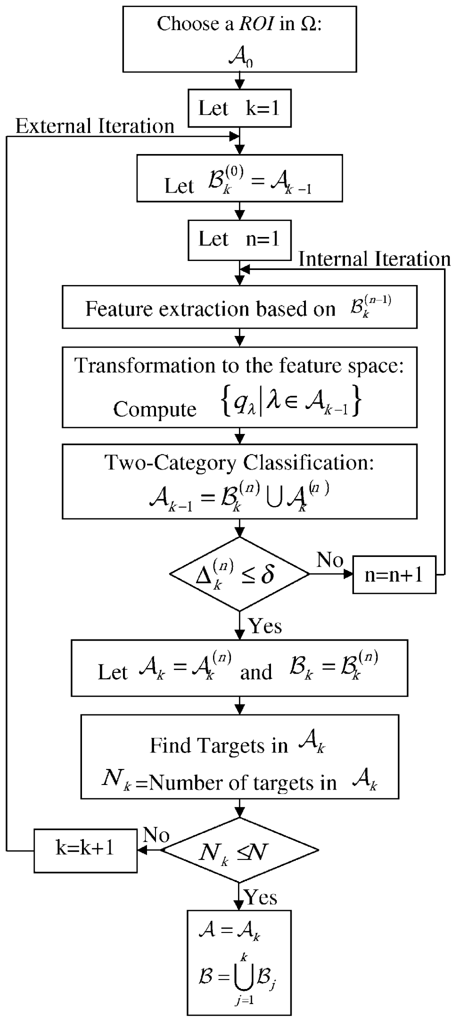

Fig. 1. Block diagram of the anomaly detection algorithm.

the partition of $\mathscr{A}_{k-1}$ into $\mathscr{A}_k$ and $\mathscr{B}_k$. This step itself is carried out in an iterative manner. A first guess of background subset, denoted by $\mathscr{B}_k^{(0)}$ is chosen in $\mathscr{A}_{k-1}$. For each pixel in $\mathscr{B}_k^{(0)}$, an observation vector, r_λ, is produced (containing the values of the pixels in the neighborhood of the pixel λ). The most dominant principal components of $\{r_\lambda | \lambda \in \mathscr{B}_k^{(0)}\}$ are chosen to be the features which span the feature space. A disjoint partition of $\mathscr{A}_{k-1}$ into $\mathscr{A}_k^{(1)}$ and $\mathscr{B}_k^{(1)}$ is then carried out using the single hypothesis scheme, as described in Section 3. The next step includes an update of the feature space using the first most dominant principal components of $\{r_\lambda | \lambda \in \mathscr{B}_k^{(1)}\}$ and disjointly partitioning $\mathscr{A}_{k-1}$ into $\mathscr{A}_k^{(2)}$ and $\mathscr{B}_k^{(2)}$ in the new feature space. This procedure iterates, with partitioning $\mathscr{A}_{k-1}$ into $\mathscr{A}_k^{(n)}$ and $\mathscr{B}_k^{(n)}$ where the feature space is obtained at each iteration by PCA done on the background subset $\mathscr{B}_k^{(n-1)}$. The stopping rule for this iterative procedure relies on the relative change, $\Delta_k^{(n)}$, between $\mathscr{B}_k^{(n-1)}$ and $\mathscr{B}_k^{(n)}$. Specifically, we compute the relative change defined by

$$\Delta_k^{(n)} = \frac{|(\mathscr{B}_k^{(n)} \cup \mathscr{B}_k^{(n-1)}) \setminus (\mathscr{B}_k^{(n)} \cap \mathscr{B}_k^{(n-1)})|}{|\mathscr{B}_k^{(n-1)}|}, \qquad (7)$$

where $|\mathscr{B}_k^{(n)}|$ is the number of pixels in the subset $\mathscr{B}_k^{(n)}$. Then we compare the relative change to a certain threshold, δ, and stop the iterations when $\Delta_k^{(n)} \leqslant \delta$. The relative change is defined here as the ratio between the number of pixels in the exclusive-or set of $\mathscr{B}_k^{(n)}$ and $\mathscr{B}_k^{(n-1)}$, and the total number of pixels in $\mathscr{B}_k^{(n-1)}$. Small relative change indicates an insignificant difference between the present background subset and the previous one.

The diagram in Fig. 1 shows also an external loop for partitioning the anomalous subsets iteratively. Its stopping rule is based on the number of the potential targets in the anomalous subset. Specifically, the number of potential targets in $\mathscr{A}_k$ should be smaller than or equal to the presumed maximal number of possible targets in the given ROI. The number of potential targets is determined without much a priori knowledge about the targets, by clustering the pixels in the spatial domain and identifying a cluster as target based on the size and shape of the cluster. Morphological tools are used for finding connected sets with areas which fit the presumed size. These are classified as potential targets and counted as such. The number of potential targets, N_k, in the anomalous subset, is then compared to the presumed maximum number of targets, N. The external iterative procedure proceeds until $N_k \leqslant N$.

5. Example

Sea-mines detection in shallow water involves addressing the varying shape of the ocean surface and vegetation [3] which yields large variability in background clutter. Conventional methods using a preprocessing procedure and matched filters [7] for detection of mine-size regions, that closely match a typical mine signature, presume a priori knowledge of the mine's size and shape. Normally, mines in a sonar image include a highlight region, representing the mine's body and a dark region representing its shadow. The variability in the background clutter and mine appearance in the sonar image, leads to a high false alarm rate in conventional methods. Fig. 2(a) presents two sonar images with sea-mines (one in each image). These images include different clutter patterns. The left image in Fig. 2(a) contains a clutter pattern of highlights and shadows whereas the particular mine in this image contains no highlight. The clutter in the right image has a different pattern containing mine-like shaped blobs. Conventional methods fail in detecting sea-mines in such varying environments. Furthermore, such clutters may produce a large number of false alarms. An experienced human viewer would be able to detect the mines based on the difference from the background clutter. The proposed algorithm mimics the detection mechanism of the human viewer by detecting anomalies in the background clutter.

The application of the proposed algorithm to the sonar images is demonstrated in Fig. 2. The algorithm employs the three most dominant principal components of 5×5 pixel neighborhoods. This neighborhood size characterizes the information which differentiates between the background texture and mines. The relative change threshold was set to 1%, the confidence level was set to 95%, the presumed maximum number of targets was set to two. The coarse target detection procedure was based on spatial morphologic detection of anomalous clusters composed of 15 pixels or more. Fig. 2(b) show the result of the first external iteration, containing several potential targets (connected pixel

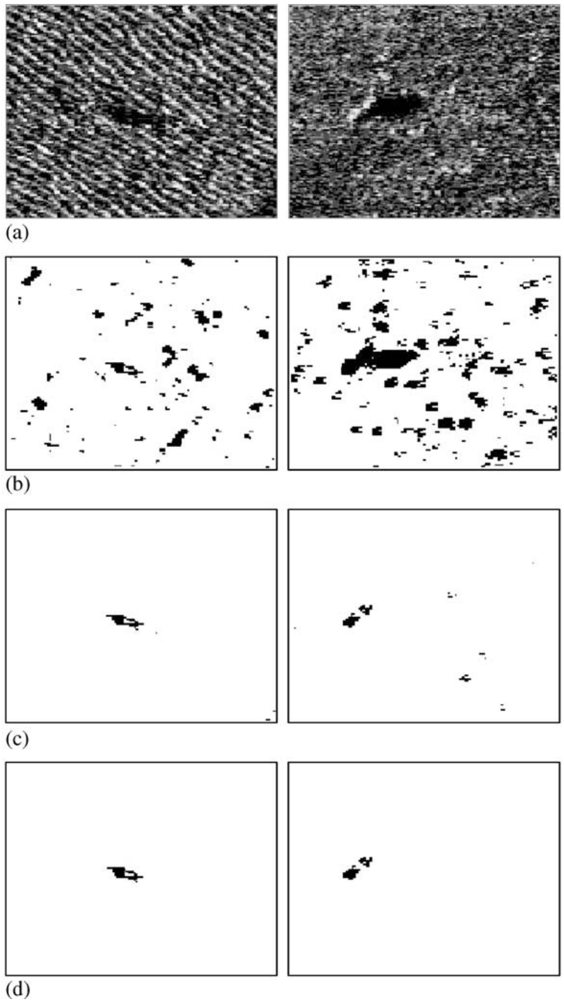

Fig. 2. (a) Original sonar images containing sea-mines. (b) The corresponding images of first iteration anomalies, $\mathscr{A}_1$ (black pixels). (c) The second iteration anomalies, $\mathscr{A}_2$. (d) The result of a morphological filtering for coarse target detection.

groups of at least 15 pixels). The second iteration, as shown in Fig. 2(c), significantly reduces the number of potential targets. These potential targets are presented in Fig. 2(d) after the morphological filtering.

6. Conclusion

The proposed algorithm includes an iterative procedure of feature extraction, based on PCA of the spatial information around each pixel. The background is statistically characterized in the feature space, and a single hypothesis scheme is used for the detection of anomalous pixels. This method was successfully employed on a large data-set of sea-mines sonar images.

The method proposed in this paper can be extended to hyper-spectral imagery [8] by incorporating the additional spectral information, related to each pixel, into the feature vectors. Rather than considering each pixel spectra separately, the spatial information is combined with the spectral data to improve the anomaly detection performance. Multi-resolution representations [4] may as well improve the detection performance, particularly in cases the background textures are nonuniform, and in cases of no a priori information regarding the sizes and shapes of targets is available.

References

[1] E.A. Ashton, Detection of subpixel anomalies in multispectral infrared imagery using an adaptive Bayesian classifier, IEEE Trans. Geosci. Remote Sensing 36 (March 1998) 506–517.

[2] G.J. Dobeck, J. Hyland, Sea mine detection and classification using side-looking sonars, in: Proceedings of the SPIE Annual International Symposium on Aerospace/Defense Sensing, Simulation and Control, Vol. 2496, 1995.

[3] G.J. Dobeck, J.C. Hyland, L. Smedley, Automated detection/classification of sea mines in sonar imagery, in: Proceedings of the SPIE, Vol. 3079, Orlando, FL, April 1997.

[4] C.H. Fosgate, H. Krim, W.W. Irving, W.C. Karl, A.S. Willsky, Multiscale segmentation and anomaly enhancement of sar imagery, IEEE Trans. Image Process. 6 (1) (January 1997) 7–20.

[5] K. Fukunaga, Introduction to Statistical Pattern Recognition, 2nd Edition, Academic Press, New York, 1990, pp. 66–77 (Chapter 3).

[6] L. van Kempen, H. Sahli, Signal processing techniques for clutter parameters estimation and clutter removal in GPR data for landmine detection, in: Proceedings of the 11th IEEE Workshop on Statistical Signal Processing, Singapore, August 2001, pp. 158–161.

[7] D. Manolakis, G. Shaw, Detection algorithms for hyperspectral imaging applications, IEEE Signal Process. Mag. (January 2002) 29–43.

[8] D.W.J. Stein, S.G. Beaven, L.E. Hoff, E.M. Winter, A.P. Schaum, A.D. Stocker, Anomaly detection from hyperspectral imagery, IEEE Signal Process. Mag. (January 2002) 58–69.

[9] R.A. Weisenseel, W.C. Karl, D.A. Castanon, C.A. Dimarzio, Statistical approach to multi-channel spatial modeling for the detection of mine-like targets, in: Proceedings of the SPIE, Detection and Remediation Technologies for Mines and Mine-Like Targets, Vol. 3392, April 1998, pp. 1044–1053.