Multi-Sonar Adaptive Target Detection using the Sphericity Test

Nick Klausner and Mahmood R. Azimi-Sadjadi

Department of Electrical and Computer Engineering
Colorado State University, Fort Collins, Colorado 80523-1373
Email: {nklausne,azimi}@engr.colostate.edu

***Abstract*—Detection of underwater objects in sonar imagery is highly challenging task. Human operators often times outperform computer algorithms in moderate and difficult environments due to the high variability in environmental conditions and human proficiency in incorporating in-situ information and context. Most automatic target recognition algorithms are pre-trained in a different environment leading to high levels of model mismatch when new environments are encountered. This paper presents an environmentally adaptive algorithm for underwater target detection in multiple sonar images using the Sphericity Test. By estimating local second-order statistics within the image, the proposed application of the Sphericity Test is capable of in-situ adaptation in changing underwater environments. The proposed detector is applied to a dataset consisting of pairs of beamformed images collected over multiple target fields. The method will also be compared and benchmarked to a similar method which does not employ environmental adaptation. Performance measures including probability of detection/false alarm and Receiver-Operator Characteristic (ROC) curve metrics will be used to evaluate the performance of each detection method.**

***Index Terms*—Sphericity Test, Generalized Likelihood Ratio Test, environmental adaptation, underwater target detection**

I. Introduction

A critical need of the U.S. Navy is the development of a reliable, efficient and robust underwater target detection system that can operate in real-time with multiple disparate sensor systems and in different environmental and operating conditions. For such a task, human operators typically outperform computer algorithms in moderate and difficult environments due to the high variability in environmental conditions and human proficiency in incorporating in-situ information and context. Most target detection and recognition algorithms in use to date are globally pre-trained across many environments and employ little or no in-situ adaptation leading to high levels of model mismatch when new environments are encountered. Thus, new methods capable of robustly adapting the parameters of the detector in-situ to the changes in the background clutter are needed.

Environmentally adaptive detection is discussed in [1] where the author identifies 5 major limitations which conventional detection algorithms suffer:

1) Many algorithms invariably assume that the image quality is uniformly excellent across the entire image.
2) Many algorithms do not exploit the range-dependent nature in sonar imagery.
3) Most methods are not tailored to the environmental conditions in which the surveying is being conducted as they rely on training and lack the ability to adapt to new environments.
4) The threshold used for detection lacks any physical meaning.
5) Most algorithms lack real-time implementation capability as they require the entire image for processing.

The author goes on to propose an adaptive algorithm that explicitly accounts for environmental characteristics, image quality, and other physics-based aspects of the problem. The algorithm is then demonstrated on a real sonar dataset collected at sea. While not specifically addressing these limitations from a physical point of view, we do somewhat address them from a statistical perspective. By relying on maximum likelihood (ML) estimation of second-order moments, the method proposed in this paper is locally adaptive making it independent of image quality, range from the sonar, and any pre-defined training set. Moreover, while not physically motivated, the threshold associated with the proposed method is chosen to statistically achieve a prescribed false alarm rate.

The detection method presented in this paper can be likened to a class of detection strategies known as anomaly detection [2], [3]. Image anomaly detection is a process of distilling a small number of clustered regions of pixels, which differ from the image's general characteristics. The type of image, its characteristics, and the type of anomalies depend on the application at hand. Common applications include detection of targets in images, detection of mine features in side-scan sonar, detection of tumorous areas in medical imaging, and detection of faults in seismic data. In [2], the authors propose modeling the wavelet coefficients in a multiscale analysis using a multidimensional Generalized Autoregressive Conditional Heteroscedasticity (GARCH) model. The proposed GARCH model is shown to generalize the causal Gauss Markov random field (GMRF) [4] model and the authors subsequently develop a matched subspace detector [5] for detecting anomalies in GARCH clutter. This work was extended in [3] by proposing the use of the noncausal autoregressive-autoregressive conditional heteroscedasticity (AR-ARCH) model. Again modeling the clutter in the image by fitting the data in the image, a matched subspace detector is constructed to perform anomaly

978-0-933957-40-4 ©2013 MTS

detection. The method is then demonstrated on a dataset of side-scan sonar images.

This paper is organized as follows. Section II develops the Sphericity Test and its environmentally adaptive application. The performance of this detection test is then demonstrated in Section III on a dataset of pairs of coregistered high frequency (HF) and broadband (BB) side-scan sonar images and compares the method to a similar approach which does not employ environmental adaptation. Finally, concluding remarks are made in Section IV.

II. Detection in Sonar Imagery using the Sphericity Test

In this section we present the Generalized Likelihood Ratio Test (GLRT) used to discriminate Regions of Interest (ROIs) containing targets from those only containing background. The GLRT [6] is an efficient and robust method of performing binary hypothesis testing when the probabilistic description of the data under one or both hypotheses depends on parameters that are assumed to be deterministic but unknown. The detection strategy taken in this paper relies on accurate estimation of the second-order statistics of background within the sonar image and the ability to determine if and when the observed second-order statistics within an ROI deviate from that of background. Those ROIs that exhibit significant deviation from background are then flagged as possibly containing a target.

For the development of the GLRT used to detect objects in sonar imagery, we assume that we are given a set of M vectors $\{\boldsymbol{x}_i\}_{i=1}^M$ with $\boldsymbol{x}_i \in C^d$ which are a set of *iid* zero mean, complex normal random vectors with covariance matrix R, i.e. $\boldsymbol{x}_i \sim \mathcal{CN}(\mathbf{0}, R)$. We also assume that the second-order statistics of the background within the sonar image is known and represented by the Hermitian, positive-definite (PD) matrix R_0. Finally, to discriminate target from background we consider the following hypothesis test

$$\begin{aligned} \mathcal{H}_0 &: \quad R = R_0 \\ \mathcal{H}_1 &: \quad R \neq R_0 \end{aligned}$$

Collecting all M *iid* copies of the random vector $\boldsymbol{x}_i$ into the data matrix $X = [\boldsymbol{x}_1 \ \boldsymbol{x}_2 \ \cdots \ \boldsymbol{x}_M] \in C^{d\times M}$, this data matrix has probability density function (PDF)

$$\begin{aligned} f(X;R) &= \prod_{i=1}^{M} f(\boldsymbol{x}_i;R) \\ &= \prod_{i=1}^{M} \frac{1}{\pi^d \det(R)} \exp\left\{-\boldsymbol{x}_i^H R^{-1}\boldsymbol{x}_i\right\} \\ &= \frac{1}{\pi^{dM}\det(R)^M} \exp\left\{-M\mathrm{tr}(R^{-1}\hat{R})\right\} \end{aligned}$$

where $\hat{R}$ is the sample covariance matrix

$$\hat{R} = \frac{1}{M} X X^H = \frac{1}{M}\sum_{i=1}^{M} \boldsymbol{x}_i \boldsymbol{x}_i^H$$

The GLRT rejects $\mathcal{H}_0$ for small values of the likelihood ratio

$$\Lambda = \frac{f(X;R_0)}{\max\limits_{R\in\mathcal{R}} \ f(X;R)}$$

with $\mathcal{R}$ denoting the space of all Hermitian, PD matrices. In other words, the value Λ represents the likelihood that the measurements were produced by the covariance R_0, normalized by that which is maximally achievable. Using the fact that the sample covariance matrix is the maximum likelihood estimator [7]

$$\hat{R} = \arg\max_{R\in\mathcal{R}} \ f(X;R)$$

and substituting this into the expression for Λ, the likelihood ratio raised to the power $\frac{1}{dM}$ becomes

$$\tilde{\Lambda} = \Lambda^{1/dM} = \frac{f(X;R_0)}{f(X;\hat{R})} = \frac{\left(\det\hat{S}\right)^{1/d}}{\frac{1}{d}\mathrm{tr}\hat{S}} \qquad (1)$$

where $\hat{S} = R_0^{-1/2}\hat{R}R_0^{-H/2}$ is a matrix sometimes referred to as the signal-to-noise (SNR) matrix [8]. If the matrix $\hat{S}$ has the orthogonal expansion $\hat{S} = \sum_{i=1}^d \lambda_i \mathbf{u}_i\mathbf{u}_i^H$ for some set of orthonormal vectors $\{\mathbf{u}_i\}_{i=1}^d$, then the test statistic can equivalently be written as the ratio of the geometric mean to the algebraic mean of the eigenvalues of $\hat{S}$

$$\tilde{\Lambda} = \frac{\left(\prod_{i=1}^{d}\lambda_i\right)^{1/d}}{\frac{1}{d}\sum_{i=1}^{d}\lambda_i} \qquad (2)$$

This classical test is typically referred to, in multivariate statistics, as the Sphericity Test [7], [9] as the contours of equal density for the Normally distributed random vector $R_0^{-1/2}\boldsymbol{x}$ are spheres when $\mathcal{H}_0$ is true.

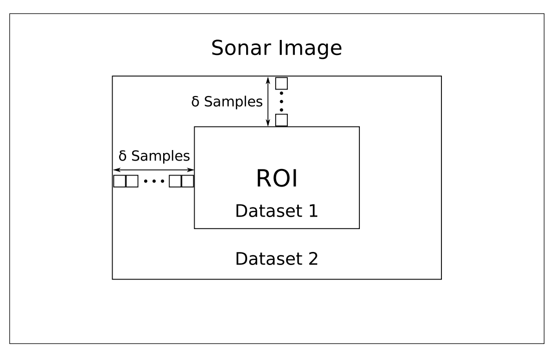

Fig. 1. Formation of Both Datasets.

As mentioned at the beginning of this section, the covariance model R_0 is interpreted to be the parameter of the Normal distribution describing realizations of background within the image. However, rather than identifying this parameter and leaving it constant throughout the entire data set (via a global pre-training process using different clutter densities for example) a modified approach to detection is taken here in which the covariance matrix R_0 is estimated from data that is spatially local to the ROI of interest by extracting the δ observations (blocks) in both range and cross-range nearest to that ROI. In this sense, the detector simply becomes one which detects local statistical anomalies within the sonar image [2], [3]. This process is depicted in Figure 1. Our detection hypothesis in this case is that, by whitening observations within the ROI using the data immediately surrounding that ROI (Dataset 2), we produce a new observation which exhibits a distribution that significantly deviates from a spherical behavior when that ROI (Dataset 1) contains a target. However, when that same ROI contains background alone we expect these whitened observations to follow a distribution which is more spherical in nature.

Looking closely at the matrix $\hat{S} = R_0^{-1/2}\hat{R}R_0^{-H/2}$ and assuming that our data is truly Gaussian distributed, we can see that under the $\mathcal{H}_0$ hypothesis this matrix is distributed according to a complex Wishart distribution with a d-dimensional identity matrix and M degrees of freedom [10]

$$M\hat{S} \mid \mathcal{H}_0 \sim \mathcal{CW}_d(I, M)$$

Thus, it becomes clear that the null distribution of the test is completely independent of the covariance matrix R_0 used to whiten the observations within the ROI. In fact, using classical theory on the asymptotic null distribution of log-likelihood ratios, one can show [9] that for large M

$$2d\left[\frac{2d^2+1}{6d} - (M-1)\right]\ln\Lambda \overset{M\to\infty}{\longrightarrow} \chi^2_{d^2-1}$$

which is clearly in no way dependent on the matrix R_0. Consequently, if the detector is implemented in the Neyman-Pearson [7] sense by choosing a threshold that achieves a particular false alarm rate, there is theoretically no reason to adapt the threshold even though we choose to adapt R_0 using the data that is spatially local to the ROI of interest. In other words, the null distribution of the test and thus the threshold itself remain invariant to the covariance model used to describe background within the image.

III. Test Results

A. Data Description and Pre-Processing

The proposed environmentally adaptive Sphericity Test is applied to a dual-channel sonar data set consisting of one HF high-resolution side-scan image as well as one BB sonar image which both cover the same region on the seafloor with the images captured from multiple target fields. Co-registration of these sonar systems is guaranteed as both HF and BB sonar are mounted on the same Autonomous Underwater Vehicle (AUV) and use the same receiver array. The pinging for the HF and BB sonar systems is done simultaneously as they are sufficiently far apart in frequency such that the returns are easily separable. Thus, the sonar images used in this study are disparate in frequency content and resolution but not in location, elevation, or aspect. Each image is complex-valued and generated at the output of the k-space or wavenumber beamformer [11]. The image database contains 61 co-registered sonar images with each image consisting of both port and starboard-side images. The database contains 77 targets with some images containing more than one target. Because the HF sonar provides higher spatial resolution and better ability to capture target details and characteristics while the BB sonar offers much better clutter suppression ability with lower spatial resolution, detectors were constructed using both the HF and BB sonar images together to ensure a high probability of detection with a low false alarm rate.

When processing the images in the data sets for the detector, each pair of HF and BB images is first partitioned into coregistered ROIs with 50% overlap in both the vertical and horizontal directions. ROIs are formed in an overlapping fashion to ensure that the target will not be split among different ROIs. Thus, if an ROI contains a target, it will encompass the entirety of the target structure. Based on the average target size, ROIs pertaining to HF images are chosen to be 84 pixels tall by 144 pixels wide. Because of differences in beamwidth in HF and BB sonar leading to discrepancies in image resolution, the ROIs pertaining to BB images are not the same size and are chosen to be 28 pixels tall by 288 pixels wide. This choice of ROI sizes ensures proper correspondence among the HF and BB images.

Once the pair of coregistered ROIs has been extracted from each of the sonar images, each ROI is partitioned into non-overlapping blocks of size 6×6 for HF images and 2×12 for BB. The difference in block size for each sonar type is again a byproduct of their difference in resolution. Corresponding blocks in the ROIs are then reshaped into vectors and concatenated to form the $d = 60$ dimensional composite observation vector $\boldsymbol{x}_i$ defined in Section II. The data matrix X is then formed from all $M = 336$ blocks $\left(\frac{84\times144}{6\times6} = \frac{28\times288}{2\times12} = 336\right)$ in each pair of coregistered ROIs. Figure 2 gives a graphical overview of the ROI formation, blocking, and ensemble pre-processing steps used to prepare the multi-channel sonar data for detection using an arbitrary N number of sonar images (e.g., one HF and many BB), although for this particular application $N = 2$. To make the detector adaptive to the statistics of the background, for every particular ROI of interest an additional δ samples (blocks) are extracted in both the along-track and cross-track dimensions as explained in Section II and displayed in Figure 1. Here we choose $\delta = 4$ in which case we extract 368 samples which are then used to estimate the background covariance matrix R_0.

B. Detection Results

The environmentally adaptive Sphericity-based Test was then implemented on the sonar imagery dataset. To get an idea of how much the adaptive version of the test improves over a situation where one does not estimate R_0 from the surrounding background, the detector was compared to a similar method that uses the exact same likelihood ratio given in (1) or (2) but uses a fixed covariance model R_0 using an

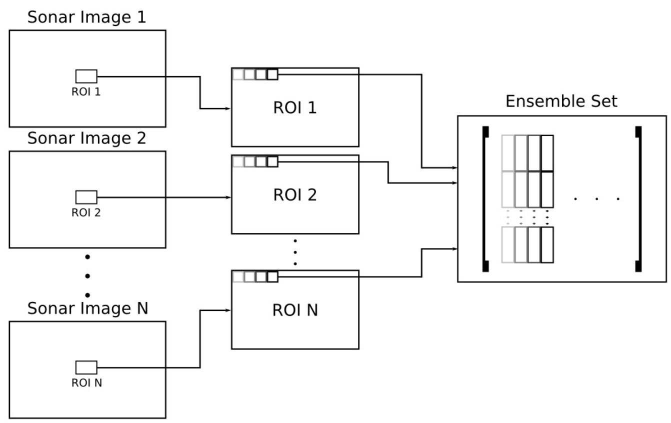

Fig. 2. Formation of the ensemble data matrix X using multiple sonar images.

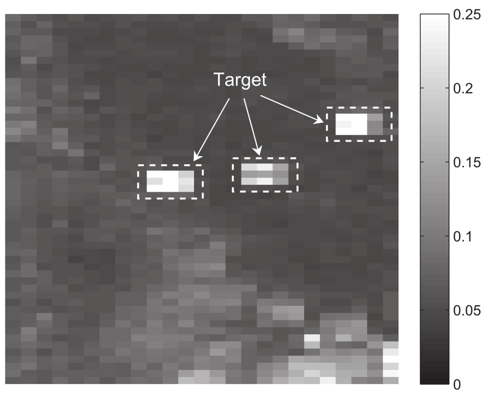

(a) Adaptive

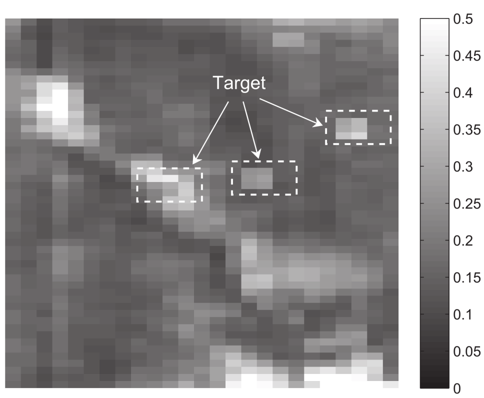

(b) Non-Adaptive

Fig. 3. Likelihood map using both the adaptive and non-adaptive detectors for one image.

appropriately chosen training set. For this test, the training set was constructed using multiple ROIs containing a range of background difficulty using a subset of 25 images from the dataset. Figure 3 displays the likelihood ratio for each ROI within one image for both (a) the method which relies on local statistics to estimate R_0 ('Adaptive') and (b) the method which trains and fixes R_0 ('Non-Adaptive'). Note that this particular likelihood map (actual image not shown) contains three targets which are shown with dashed boxes in Figure 3. From these images, we can see that the both methods produce large likelihood for the three targets in this image. However, it is also clear that the adaptive version of the test remains more robust to the changing background conditions within the image as the contrast between target and background is much greater. Also, one can see that the adaptive version does not produce a large number of false detections compared to its non-adaptive counterpart.

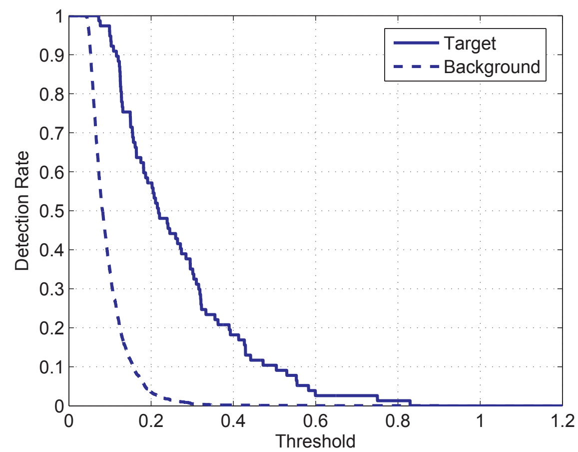

(a) Adaptive

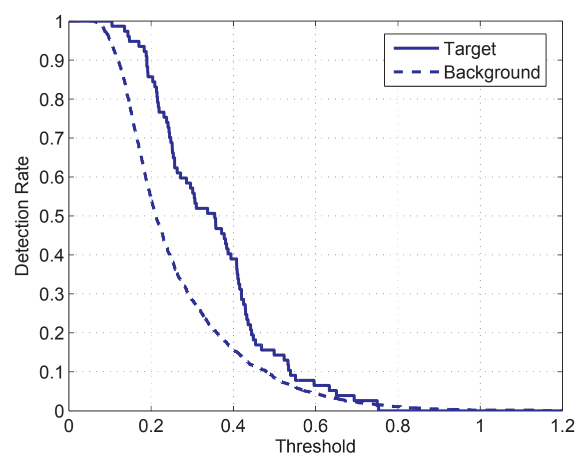

(b) Non-Adaptive

Fig. 4. Detection rates for both detection methods.

IV. Conclusions

In this paper, we consider the use of the Sphericity Test to determine locations within sonar images which exhibit second-order statistics that significantly deviate from that of surrounding background for the detection of underwater

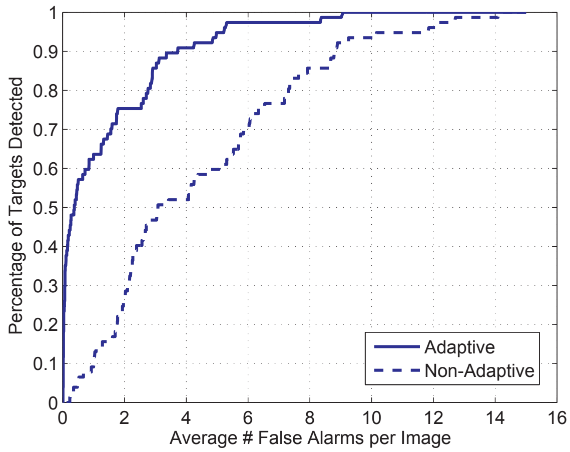

Fig. 5. ROC curves for both detection methods.

targets. In its general form, the Sphericity Test is a GLRT that tests whether or not an estimated covariance matrix is equal to a given covariance matrix using the arithmetic and geometric averages of the eigenvalues of an estimated SNR matrix. Rather than fixing the statistical model for background, a modified approach is proposed which estimates this second-order information using the data spatially local to the region of interest. Under the null hypothesis, the distribution of the statistic is independent of this statistical model making the threshold invariant to its adaptation. Both the proposed environmentally adaptive version of the Sphericity Test and its non-adaptive counterpart found by estimating the background covariance using a training dataset were then implemented and compared using a dataset of coregistered pairs of HF and BB sonar images containing underwater targets. Throughout this study, it was not only shown that the adaptive Sphericity Test performs well at discriminating target from background but also that it by far outperforms any similar method based on training.

Both the adaptive and non-adaptive versions of the Sphericity Test were then applied to all of the images in the dataset. Figures 4 (a) and (b) display right tail probabilities, $P[\Lambda > \lambda]$, as a function of the threshold λ for the adaptive and non-adaptive detectors, respectively, and for all 77 targets in the dataset (shown by a solid line) along with background (shown with a dashed line). In other words, for a given value of λ along the x-axis of these plots, each curve gives the percentage of ROIs containing either background or target that fall above that threshold. From these two plots, we can see that the adaptive version of the test is sufficiently capable of discriminating target from background while the separation between target and background likelihood values is smaller for the non-adaptive version. Finally, Figure 5 displays a ROC curve showing the percentage of targets detected versus the average number of false alarms per image for both detection methods. Again, from this figure we can see that the adaptive version of the test outperforms its non-adaptive counterpart. At an average of 4 false alarms per image, the adaptive method detects over 90% of the targets in this dataset while the non-adaptive version detects little over half of the targets at the same false alarm rate.

ACKNOWLEDGMENT

This work was supported by the Office of Naval Research, Code 321OE under contract #N00014-12-C-0017.

REFERENCES

[1] D. P. Williams and J. Goen, "A fast physics-based, environmentally adaptive underwater object detection algorithm," *IEEE Oceans, Santander, Spain*, pp. 1–7, 2011.

[2] A. Noiboar and I. Cohen, "Anomaly detection based on wavelet domain GARCH random field modeling," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 5, pp. 1361–1373, May 2007.

[3] S. Mousazadeh and I. Cohen, "Anomaly detection in sonar images based on wavelet domain noncausal AR-ARCH random field modeling," *2010 IEEE Convention of Electrical and Electronics Engineers in Israel*, pp. 306–309, 2010.

[4] J. W. Woods, "Two-dimensional discrete Markovian fields," *IEEE Transactions on Information Theory*, vol. IT-18, no. 2, pp. 232–240, March 1972.

[5] L. L. Scharf and B. Friedlander, "Matched subspace detectors," *IEEE Transactions on Signal Processing*, vol. 42, no. 8, pp. 2146–2157, August 1994.

[6] S. M. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*. Prentice Hall, New Jersey, 1998, vol. 2.

[7] R. Muirhead, *Aspects of Multivariate Statistical Theory*. John Wiley and Sons, New Jersey, 2005.

[8] L. L. Scharf and B. D. Van Veen, "Low rank detectors for Gaussian random vectors," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 35, no. 11, pp. 1579–1582, Nov 1987.

[9] D. B. Williams and D. H. Johnson, "Using the sphericity test for source detection with narrow-band passive arrays," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 38, no. 11, pp. 2008–2014, November 1990.

[10] N. R. Goodman, "Statistical analysis based on a certain multivariate complex Gaussian distribution (An introduction)," *Annals of Mathematical Statistics*, vol. 34, no. 1, pp. 152–177, March 1963.

[11] M. Soumekh, *Synthetic Aperture Radar Signal Processing*. Wiley, 1999.