FINAL REPORT

UXO Detector for Underwater Surveys Using Low-Frequency Sonar

SERDP Project MR-2415

JUNE 2015

Robert van Vossen
TNO

Alan Hunter
University of Bath

Distribution Statement A
This document has been cleared for public release

SERDP
DOD ▪ EPA ▪ DOE

This report was prepared under contract to the Department of Defense Strategic Environmental Research and Development Program (SERDP). The publication of this report does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official policy or position of the Department of Defense. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the Department of Defense.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE *(DD-MM-YYYY)*
01-06-2015

2. REPORT TYPE
Technical Report

3. DATES COVERED *(From - To)*
01-04-2014 - 01-06-2015

4. TITLE AND SUBTITLE
UXO Detector for Underwater Surveys using Low-Frequency Sonar

5a. CONTRACT NUMBER
W912HQ-14-C-005

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
A.J. Duijster, A.J. Hunter, R. van Vossen, A.L.D. Beckers

5d. PROJECT NUMBER
MR-2415

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

TNO
Oude Waalsdorperweg 63
The Hague
The Netherlands

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

SERDP
4800 Mark Center Drive,
Suite 17D08
Alexandria VA22350-3600

10. SPONSOR/MONITOR'S ACRONYM(S)

SERDP

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release: distribution is unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Locating and surveying underwater dumping sites of unexploded ordnance (UXO) is challenging, particularly when they are buried beneath the seafloor sediment. Low-frequency (LF) sonars have a demonstrated capability for detecting buried objects. However, high clutter densities observed in relevant operational environments will limit their use unless a capacity is developed to distinguish between clutter (e.g., natural objects, debris, and seafloor topography) and UXO. In this project, a two-stage detection approach has been developed for UXO detection. This approach has been applied to MUD-2011 data and the evaluation demonstrated that it is capable of reducing the number of contacts corresponding to natural clutter with 60% compared to a threshold detector.

15. SUBJECT TERMS
Sonar, Unexploded Ordnance (UXO), Detection

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			R. van Vossen
U	U	U	UU	39	**19b. TELEPHONE NUMBER** *(include area code)* +31 8886 61052

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Contents

1 Abstract 1
2 Introduction 3
2.1 Background 3
2.2 Outline of Report 4
3 MUD System and 2011 Trial 5
4 Existing Processing Chain (MR-2200) 6
4.1 Vertical Beamsteering 6
4.2 Synthetic Aperture Processing 6
4.3 Target / Background Separation by Incoherent Wavelet Shrinkage 6
5 Detection Chain 7
5.1 Contact Detection 8
5.2 Signature Extraction 9
5.3 Feature Extraction 9
5.3.1 Image-based Features 10
5.3.2 Multi-Aspect Acoustic Colour Features 10
5.4 Clutter Reduction by Contact Classification 13
6 Results 15
6.1 Contact Detection 16
6.2 Extracted signatures 19
6.3 Clutter suppression based on image features: Preliminary results 20
6.4 Clutter suppression based on image features: Classification results 22
7 Discussion 23
8 Conclusions 24
9 Way ahead 24
10 Publications 25
References 25
Appendix A - D Addendum 27
Appendix E – UA Paper 28
1. Introduction 29
2. MUD Sonar and Processing Chain 29
3. Target and Clutter Detection 30
4. Results 34

5. Conclusions and Discussion 34
6. Acknowledgements 35

List of Figures

Figure 2-1 – The complete flow diagram, showing both the processing chain and the detection chain. The numbers represent the corresponding sections in this report.

Figure 3-1 – MUD system: (a) wet end; (b) deployment.

Figure 5-1 – Classification tree: Schematic illustration of a detection and classification approach.

Figure 5-2 – Flow diagram of the contact detection and signature/feature extraction chain.

Figure 5-3 – Example illustration of contact and highlight masks.

Figure 5-4 – (a) SAS image snippets and multi-aspect acoustic colour plots from multiple runs of different geometries on the EVA cylinder. The image snippets show an area of 5×4 m around the target, and the acoustic colour plots show dependencies in aspect angle from -45° to 45° relative to the across-track direction and frequencies from 4 to 9 kHz; (b) Track orientations with respect to the target. The colours of the arrows in (a) correspond to the colour of the tracks in (b). The cross indicates the target position.

Figure 5-5 – Three ways to increasing the angular coverage of (a) the current multi-aspect acoustic colour plot: (b) by changing the hardware to obtain a larger angular coverage of a single run, (c) by combining several runs into a single plot, and (d) by using circular SAS.

Figure 5-6 – An example of Parzen classification in the training phase, with varying kernel sizes for the two classes. In the title of the subfigures, the kernel sizes are shown between brackets for the two classes respectively. Circles (○) are class 1 samples, while crosses (×) are class 2 samples. Green colour means correct classification, while red means incorrect classification.

Figure 6-1 – The extended classification tree, with the four categories A-D: UXO, deployed clutter, non-deployed clutter and background, respectively.

Figure 6-2 – The SAS image used as input for the detection process, where the thick dashed circles denote deployed UXO objects and the thin dotted circles denote deployed non-UXO objects. The inset shows the area marked with a gray square. Background removal has not been applied to this SAS image.

Figure 6-3 – The final result, showing the individual contact clusters, each containing one or more highlights. The inset shows the area marked with a gray square.

Figure 6-4 – (a) the probability of detection and (b) the relative number of false alarms, both as a function of the wavelet threshold and the detection threshold for the coherent images of 5 runs combined. The dotted line shows the default wavelet threshold value used in [5]. In the blank area, the low detection threshold leads to unreliable results due to fusion of multiple contacts into large areas.

Figure 6-5 – (a) Intensity of all contacts plotted versus intensity variation feature. Contacts below the dashed line can be identified as clutter. (b) Cluster area feature of all contacts versus intensity variation feature. Contacts outside the dashed box can be marked as clutter contacts.

Figure 6-6 – The contacts identified empirically as clutter are marked with a red colour, while the remaining contacts (targets and clutter) are marked with a green colour. The thick dashed circles denote the positions of deployed UXO objects and the thin dotted circles denote the

positions of deployed non-UXO objects. The insets show an enlarged part of the images, marked by a gray square.

Figure 6-7 – The contacts identified by Parzen classification as clutter are marked with a red colour, while the remaining contacts (targets and clutter) are marked with a green colour. The thick dashed circles denote the positions of deployed UXO objects and the thin dotted circles denote the positions of deployed non-UXO objects. The insets show an enlarged part of the images, marked by a gray square.

Figure E-1 – (a) SAS image and ground truth target positions, marked with a square; (b) the same image, after wavelet shrinkage, showing the coherent part.

Figure E-2 – Flow diagram of the image processing and target detection chain.

Figure E-3 – The binary contact clusters (after the morphological opening and closing step) for both the SAS image (a) and the contacts image (c), and the resulting detections overlayed on the SAS image (b) and the coherent image (d). Ground truth positions are marked with a cyan square, while contact detections are marked with a green circle. The colour scale is identical in the images and is normalised between 0 and 1.

Figure E-4 – The ROC curves for the SAS and post-processed images for two separate runs.

Acronyms

APL	Applied Physics Laboratory
BOSS	Buried Object Scanning Sonar
CMRE	Centre for Maritime Research and Experimentation
EDF	Empirical Distribution Functions
FA	False Alarm (Rate)
HF	High Frequency
k-NN	k-Nearest Neighbours
LF	Low Frequency
MUD	Mine Underwater Detection
NRL	Naval Research Laboratory
NSWC-PC	Naval Surface Warfare Centre, Panama City
PD	Probability of Detection
REMUS	Remote Environmental Monitoring UnitS
RNLN	Royal Netherlands Navy
ROC	Receiver Operating Characteristics
RTK-GPS	Real Time Kinematic Global Positioning System
SAS	Synthetic Aperture Sonar
SERDP	Strategic Environmental Research and Development Program
SRR	Signal-to-Reverberation Ratio
TNO	Netherlands Organisation for Applied Scientific Research
UXO	Unexploded Ordnance

Keywords

Sonar, Synthetic Aperture Sonar, Interferometry, Wavelets, Classification, Clutter reduction

Acknowledgements

This work was performed under direct support from the SERDP Program Office. We would like to acknowledge the very positive and encouraging program management role played by Dr. Herb Nelson, SERDP Program Manager for Munitions Response. The authors acknowledge the Royal Netherlands Navy for funding and operational support for the MUD sea trials.

1 Abstract

Objectives:
Locating and surveying underwater dumping sites of unexploded ordnance (UXO) is challenging, particularly when they are buried beneath the seafloor sediment. Low-frequency (LF) sonars have a demonstrated capability for detecting buried objects. However, high clutter densities observed in relevant operational environments will limit their use unless a capacity is developed to distinguish between clutter (e.g., natural objects, debris, and seafloor topography) and UXO. Simple level-based detectors generate an impracticable number of false alarms and are therefore inadequate for this purpose. The primary objective of this project therefore is to improve the LF sonar capability of detecting UXO by focusing on the procedure of reducing the number of contacts corresponding to clutter.

Technical Approach:
The processing chain developed in SERDP project MR-2200 is used as input to this project, together with data acquired by an experimental low-frequency side-looking sonar in the MUD-2011 experiment funded by the Netherlands Ministry of Defense; MUD-2011 was conducted in a shallow-water estuary with a muddy bottom and the deployed objects included 155mm shells, Mk82 and Mk84 bombs, and mine-like targets; the objects were buried in approximately 0.5-1.0 m of mud. A two-stage detection approach has been developed that uses LF SAS images as input:

1. An algorithm is developed for detecting contacts that correspond to deployed objects *and* clutter.
2. A feature-based approach is subsequently developed to reduce the amount of contacts corresponding to clutter.

A feature-based approach to reduce the amount of clutter contacts requires a careful feature selection. Feature selection therefore received substantial effort in the current project. The robustness of the feature set and the performance of the two-stage detection approach are illustrated on data acquired in the MUD-2011 experiment. Furthermore, the influence of MR-2200 background suppression on the detection performance is investigated.

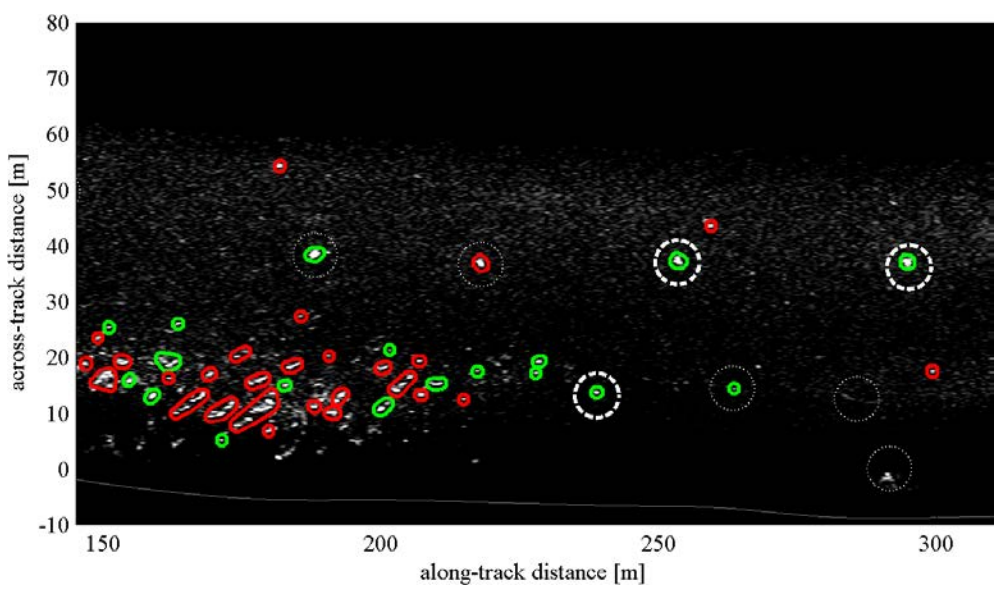

Results:

The two-stage detector has been trained and tested on different runs of the MUD-2011 experiment. The application of the two-stage detector approach has demonstrated its capability to reduce clutter. The figure illustrates the result of the detection procedure on a typical LF-SAS image. All circles correspond to detections, the red ones to contacts detected as clutter, the green ones to contacts detected as UXO; white thick dashed circles denote deployed buried UXO objects, thin dotted circles correspond to deployed non-UXO objects.

Benefits:

The UXO detector developed in this project will contribute towards achieving an operational capacity for performing rapid wide-area surveys of underwater military munitions sites using LF broadband sonar. High clutter densities are expected in relevant environments and, by reducing the false alarm rate, the burden on follow-on surveys and remediation efforts will be reduced, leading to more efficient and cost-effective operations. The output can also aid in localizing and sizing the sites and for performing check surveys after remediation. Furthermore, the results of this project will provide insights into the capabilities of low-frequency sonar and the expected performance in operational conditions.

2 Introduction

2.1 Background

High-frequency (HF) side-looking sonar (i.e., with frequencies higher than 100 kHz) is ideally suited to providing high-resolution images of the seabed. However, since sound does not penetrate into the seabed at these frequencies, such systems cannot be used for the detection of buried objects, such as unexploded ordinances (UXO). Low frequency (LF) side-looking sonar is a promising technology for the detection of objects buried in soft seafloor sediment. Acoustic energy is attenuated less by the sediment at lower frequencies and can therefore penetrate deeper, facilitating the detection of buried objects. Furthermore, a side-looking configuration yields a much higher area coverage rate compared to downward-looking systems (e.g., the BOSS system [1]) and this enables efficient surveys.

TNO has developed a LF (1 kHz – 26 kHz) side-looking sonar for experimentation on buried object detection and, with funding and support from the Dutch Ministry of Defence, have conducted sea trials in relevant environments and conditions. The system is designed to have a capability to mitigate sea surface multipath reverberation in shallow waters and to address poor directivity at low frequencies. This is achieved by the use of vertical array beamforming and synthetic aperture processing. Experimental results from the MUD-2009 and MUD-2011 sea trials demonstrate that objects buried in mud can be detected in data acquired by the LF side-looking sonar system [2],[3].

In practice a fundamental issue is that, in addition to the targets of interest (e.g., UXO), the system also observes *clutter* contacts, including other buried objects (e.g., boulders) and geological features below the mud (e.g., sand ripples). Therefore, a solution needs to be found for classifying the detections in order to discriminate between targets and clutter and thus suppress false alarms. This step is essential for the realization of a UXO detection capability.

The LF and HF classification problems are fundamentally different. While information on size and shape derived from high-resolution images are commonly used for HF classification, it is less reliable for LF classification since the wavelength is on the same order of magnitude as the dimensions of the objects of interest. In the LF frequency range, however, structural resonances may be generated in the objects. It has been indicated in experiments in controlled conditions and by modelling conducted by APL, NSWC-PC, and NRL, that useful information on the objects can be retrieved from these structural resonant features [4],[5].

The objective of this project is to explore techniques for improving the automated detection of UXO buried in mud using LF sidescan sonar in an operationally relevant environment.

2.2 Outline of Report

As input to the detection chain, data have to be acquired and processed. This is schematically illustrated in Figure 2-1. The data acquisition is summarised in Chapter 3, and the processing in Chapter 4. The main processing steps that been developed in MR-2200 are:

- Vertical beamsteering to mitigate multipath reverberation,
- SAS imaging to enhance along-track resolution, and
- Background removal to improve the signal to reverberation ratio.

In Chapter 5, the detection chain to be developed within SERDP MR-2415 is described. The approach consists of four steps:

1) Contact detection: contacts include both man-made objects and clutter,
2) Signature extraction: extract relevant contact data from a SAS image,
3) Feature extraction: compute relevant features from the extracted signatures to enable the discrimination between UXO and clutter, and
4) Contact classification: use features for clutter reduction.

Results on t he detector performance, extracted signatures, on extracted features, and on clutter reduction are presented in Chapter 6. The conclusions and the way ahead are described in Chapter 7.

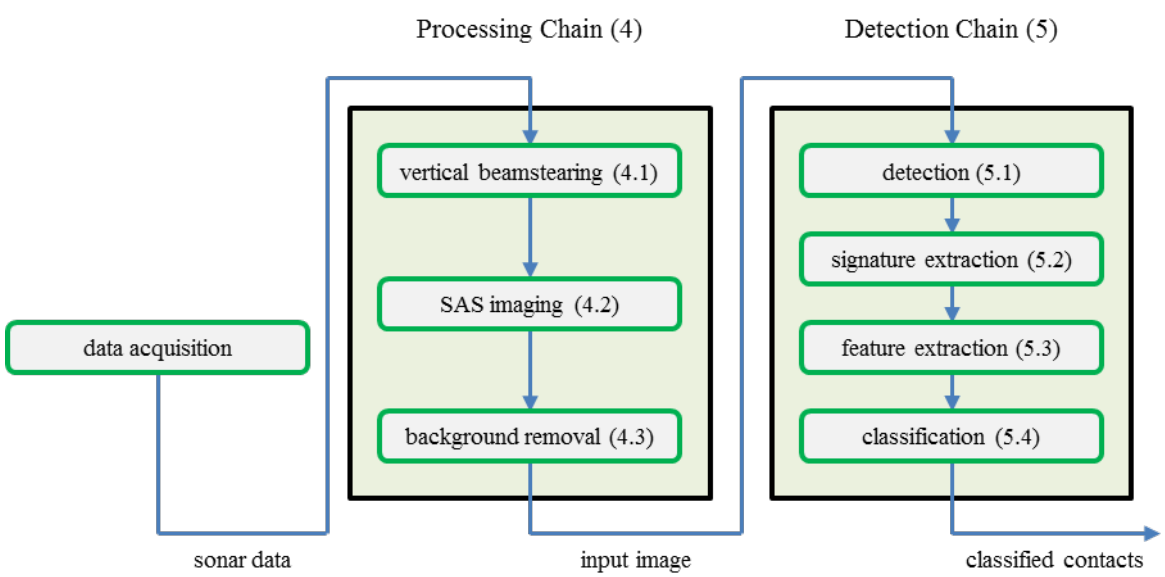

Figure 2-1 – The complete flow diagram, showing both the processing chain and the detection chain. The numbers represent the corresponding sections in this report.

3 MUD System and 2011 Trial

The wet-end of the LF sonar system is shown in Figure 3-1; it is comprised of an (exchangeable) acoustic source and two receiving arrays. One array is orientated vertically and the other horizontally, and each array is composed of 16 hydrophones. Three sources were available, covering the bandwidths from 1 kHz – 4 kHz, 4 kHz – 9 kHz, and 11 kHz – 26 kHz. The components are mounted on a frame that is adapted for operation from diver support vessels of the RNLN.

Two navigation sensors are used to monitor the position and orientation of the sonar system; these are located on top of the support frame. A photonic inertial navigation system records the 3-D accelerations and rotation angles of the system, and a real time kinematic global positioning system (RTK-GPS) provides centimeter positioning accuracy. These non- acoustical systems are necessary in order to derive the exact position and orientation of the system with respect to the test area. This accurate navigation is also required for more advanced signal processing, and in particular for synthetic aperture sonar (SAS) processing.

For the population of the test garden, a selection was made from a range of objects with different characteristics including man-made and natural objects. These targets were distributed over three lines in the test garden, with an average distance between targets of approximately 25 m. One of the deployed targets was the EVA cylinder from the Centre for Maritime Research and Experimentation (CMRE) [6]. The test garden was deployed 6 months before the trial.

The MUD-2011 trial took place from 18 April 2011 to 22 April 2011. In this period, a total of 220 runs were executed covering different parts of the test garden. REMUS control runs were conducted in February 2011 and during the trial in April 2011.

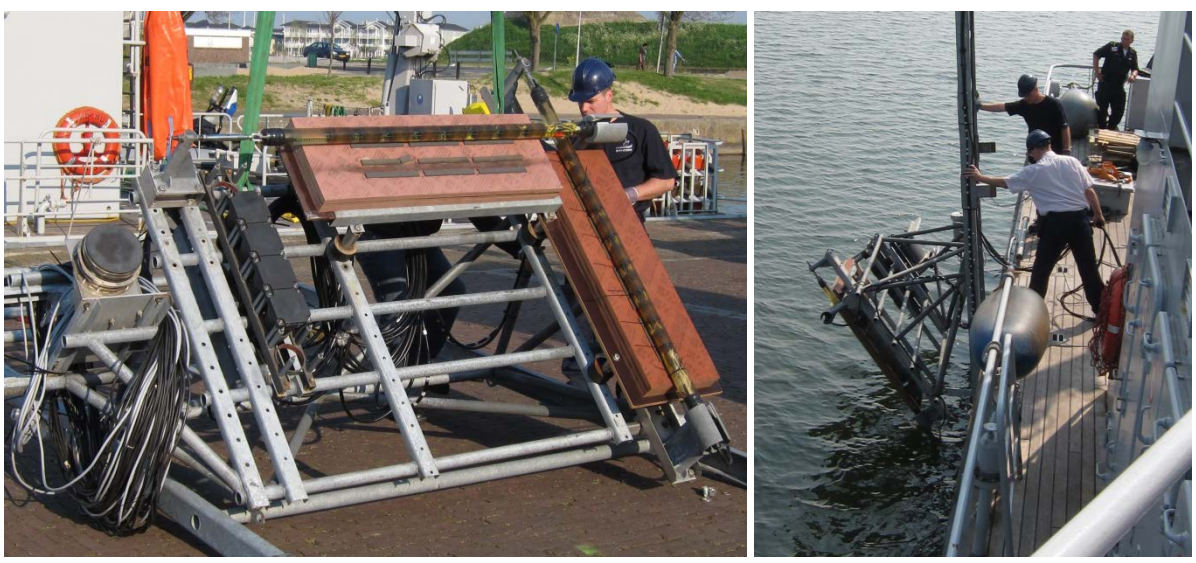

(a) (b)

Figure 3-1 – MUD system: (a) wet end; (b) deployment.

4 Existing Processing Chain (MR-2200)

The objective of the processing is to prepare the received data for optimal extraction of information for classification. The focus is on processing for the retrieval of multi-aspect acoustic colour information and on interferometric processing. In the low-frequency range considered here, the acoustic colour contains information that could be used for classification [5], and interferometric processing is used to obtain information on the seafloor bathymetry and height of objects, which may provide additional useful information on the objects and on the environment.

Generally, the processing of the received data is not straightforward due to multipath propagation in shallow water and deviations from an assumed nominal trajectory [7]. Therefore, the processing needs to be capable of dealing with these issues in order to robustly retrieve information on the targets. The current layout of the processing chain [8],[9] is outlined in Figure 2-1 and the 3 stages of the chain are described briefly in the following sections.

4.1 Vertical Beamsteering

The array of 16 vertical elements is beamsteered towards the seafloor to suppress multipath reverberation from the sea surface. This is an important processing step that enhances the quality of the image. It mitigates the receiver-side multipath contribution which can be on the same order of magnitude as the response directly obtained from the seabed. Furthermore, it is necessary for retrieving acoustic colour information on the targets, since the interference between the direct path and the path arriving via the sea-surface introduces artefacts in the acoustic colour. With vertical array beamforming, these contributions can be separated.

4.2 Synthetic Aperture Processing

The objectives of synthetic aperture processing are to enhance the signal-to-reverberation ratio (SRR) and the resolution. The enhanced SRR provided by SAS significantly improves detection performance. Furthermore, the improved resolution enables the isolation of contacts from other nearby contacts and the background reverberation.

4.3 Target / Background Separation by Incoherent Wavelet Shrinkage

Effective detection of targets in background reverberation noise requires a robust method, which is able to discern even weakly scattering objects in a highly reverberant seafloor. In order to separate the targets from the background, a coherence metric can be used. In [10], a coherence metric was derived to determine the similarity of wavelet coefficients between independent looks, i.e. different images of the same scene with statistically independent noise realizations. It is assumed that a high coherence corresponds with the reverberation-free measurements of targets, while the low coherence contributions are assumed to correspond to reverberation (background). By weighting the image according to this coherence, a separation can be made between targets and background. A thorough description of this method can be found in [10].

5 Detection Chain

The detection chain uses as input the SAS images obtained after the SAS processing described in Section 4.2 or the SAS images obtained after background removal (Section 4.3). Generally, the detection of buried targets in sonar images is difficult due to the presence of clutter and reverberation. The objective of the background removal is to suppress the reverberation. However, clutter is not suppressed by the background removal approach.

Clutter and reverberation may generate an excessive number of contacts, and a procedure has to be developed to reduce the number of contacts that do not correspond to UXO. The proposed procedure to achieve this, consists of four building blocks (Figure 2-1):

- Contact detection
- Signature extraction
- Feature extraction
- Contact classification

Contact detection aims to separate objects of interest from the background. We consider objects of interest to include deployed and non-deployed objects, where both can be natural objects or man-made objects. We further separate man-made objects into UXO and other man-made objects. The associated tree of classes is illustrated in Figure 5-1.

Subsequent steps are used to reduce the number of contacts by separating the UXO contacts from the others. Signature extraction extracts the relevant SAS image snippet as well as the associated multi-aspect acoustic colour. Feature extraction computes the features, which are used as input for the classifier. Each of these blocks is described in the subsequent sections.

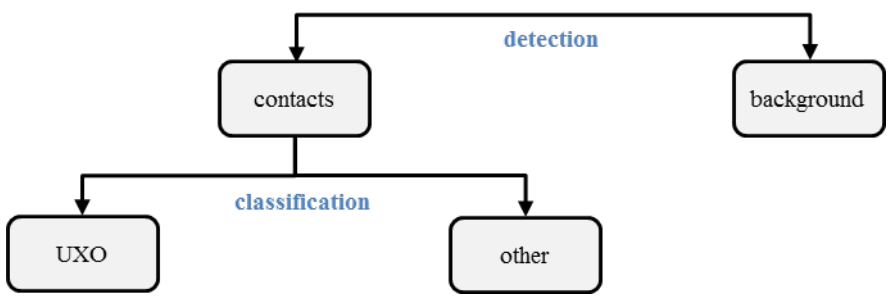

Figure 5-1 – Classification tree: Schematic illustration of a detection and classification approach.

5.1 Contact Detection

An automated approach has been developed to separate objects of interest, including UXO and clutter contacts, from the background. Once these contacts have been obtained, more detailed information can be derived using dedicated processing, for example based on image features or features of the multi-aspect acoustic colour [8].

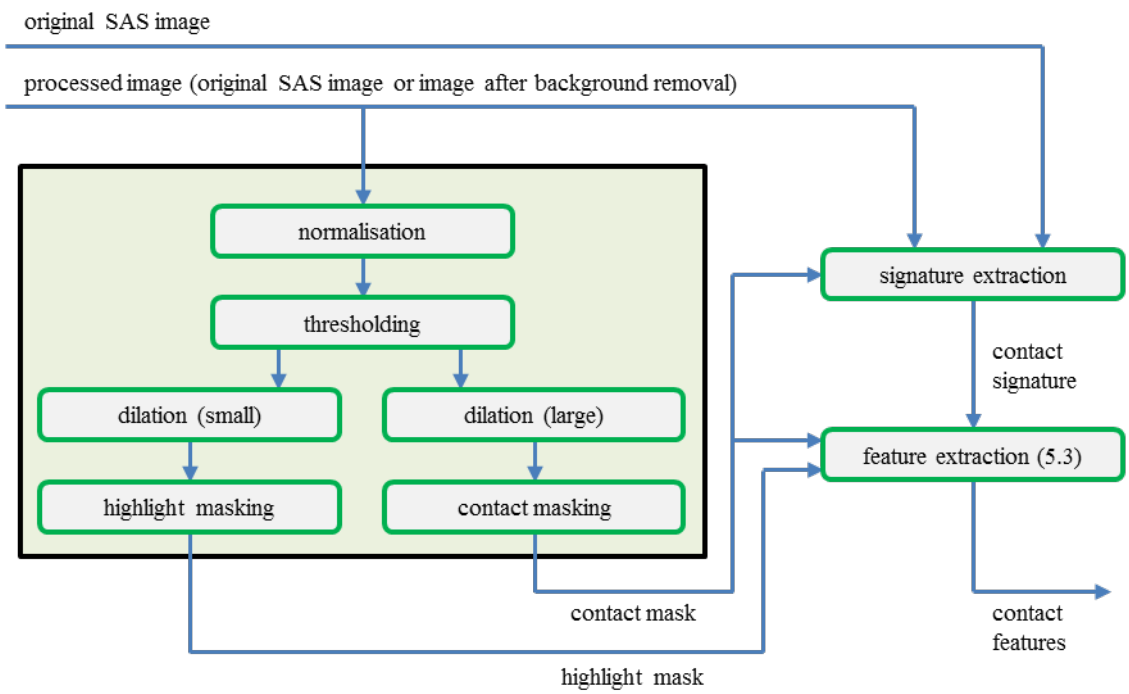

Figure 5-2 – Flow diagram of the contact detection and signature/feature extraction chain.

A flow diagram for the contact detection is given in Figure 5-2. It uses SAS images obtained after processing described in Sections 4.2 or 4.3 as input and consists of the following steps:

1) Input image normalisation.
2) Thresholding: Apply a fixed threshold to the normalised image leading to a binary image in which the objects of interest are separated from the background.
3) Dilation (small): Apply a small morphological dilation (for example, using a circular structuring element with a diameter of 50 centimeter) to the binary image in order to combine isolated groups of pixels deemed to be corresponding to the same highlight.
4) Highlight masking: generate and store masks corresponding to the individual highlights; these are used for the feature extraction (Figure 5-3).

5) Dilation (large): Apply a large morphological dilation (for example, using a circular structuring element with a diameter of 2 meter) to the binary image in order to group clusters of highlights. Each cluster represents a single contact.
6) Contact masking: Generate a list of contacts, each containing a single highlight or a cluster of highlights. Contact masks are used both in signature extraction and for feature extraction (Figure 5-3).

5.2 Signature Extraction

For each contact, the acoustic signature is extracted. This is achieved by the application of each contact mask obtained in Section 5.1 to the input SAS image. The highlight mask that corresponds to each contact is stored as well, since it is used as input for the feature extraction, described in the next section.

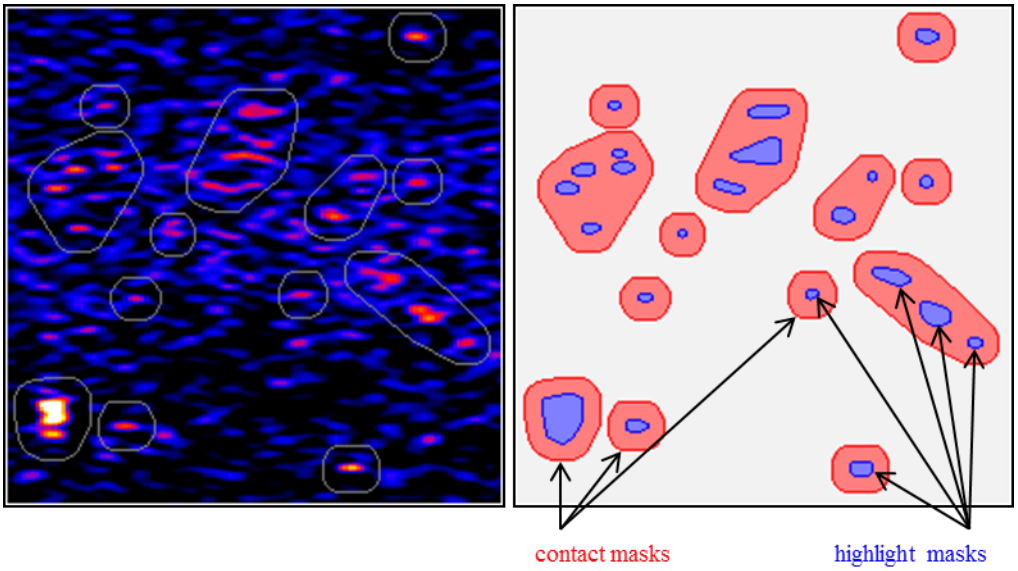

Figure 5-3 – Example illustration of contact and highlight masks.

5.3 Feature Extraction

Feature extraction calculates candidate features to be used for automated classification, here defined as the separation of UXO contacts from other contacts. Feature extraction uses both the acoustic signature (SAS snippet and acoustic colour) and the corresponding highlight mask as input.

Candidate features for object classification are generally indicators on geometrical, textural, or material properties. Given a SAS image, image processing can provide features related to size, shape, and texture of contacts. These features are referred to as *image-based features*. It is also possible to consider scattering properties of an object. This is of interest, because resonances can

be excited in the low-frequency range. The characteristics of resonances can provide information on size, shape, texture, and material properties of an object. Scattering properties of an object including resonances can be visualised in multi-aspect acoustic colour images. We explore whether features for discriminating UXO from clutter can be derived from multi-aspect acoustic colour.

5.3.1 Image-based Features

Image-based features naturally provide information on size, shape, and texture of contacts. Although it is not possible to directly obtain accurate information on the size and shape of objects in the LF range, when the wavelength is on the same order of magnitude as the objects of interest, it is anticipated that relevant information can still be derived from the image. Observations reveal that large natural bottom features (e.g. ripples) or man-made elongated objects (e.g. cables) can lead to multiple contacts within a single cluster, whereas compact features or objects (e.g. rocks, UXO) only show a single or a few highlights. Secondly, it is observed that the texture (intensity variations within a contact) of contacts corresponding to bottom features, such as ripples, often differs from the texture of man-made objects.

These observations can be used to select candidate features that aim to distinguish UXO from other contacts, such as:

- Cluster size features
- Cluster shape features
- Intensity features
- Intensity variation features

Most of the candidate features are range-independent and aspect-independent, which is beneficial because invariance generally enhances the discrimination potential.

5.3.2 Multi-Aspect Acoustic Colour Features

Multi-aspect acoustic colour is considered because this representation provides information on the angular and frequency dependent scattering properties. It contains information on the scattering strength versus angle and frequency. It is anticipated that multi-aspect acoustic colour is especially relevant for LF data. In this frequency range, the scale of objects of interest is on the same order of magnitude as the wavelength. It is thus the frequency range in which resonances could occur.

To obtain information on the scattering properties of a contact, including resonances, the SAS image is transformed to the multi-aspect acoustic colour domain using the following approach, as described in [11]:

- The SAS image snippet is used as input,
- A Fourier transform is performed, followed by a Stolt mapping, and

- A coordinate transform from wavenumber to angle is performed to obtain the acoustic colour versus angle.

Examples of multi-aspect acoustic colour images on the EVA cylinder are given in Figure 5-4, and on other objects in Appendix A. For the MUD system, these roughly cover the angle range between -45 and 45 degrees with respect to the across-track direction, and frequencies between 4-9 kHz for the runs that are considered here.

In our attempts to define candidate multi-aspect acoustic colour features, the following challenges were encountered:

- Acoustic colour plots do not only provide information on the scattering of objects, but also on the environment. Burial depth, for example, influences the frequency content of acoustic colour plots as well. For an increased burial depth of the same object, higher frequencies will be more strongly attenuated. Variations in the target response can thus also be introduced by the environment, causing a larger spread in feature values.
- Multi-aspect scattering can depend heavily on the aspect angle (e.g., Figure 5-4, Appendix A). Since the orientation of the object is generally unknown, it is difficult to exploit this information when observing only a limited range of angles. Features derived from multi-aspect acoustic colour over a limited angular range can show large variations depending on the geometry, which makes them difficult to use in a classifier. Consider for example Figure 5-4, where the acoustic colour observed for end-on target illumination is significantly different near broad side.
- Acoustic colour will also be range-dependent, because the angle of incidence will vary with range. This is especially relevant when objects are buried in sand, with a sound speed ratio between sediment and sea water larger than one. Then, penetration of sound into the seabed can be severely limited at longer ranges. The scattering properties of the object could show some variability as well with incidence angle.

We expect, however, that the angular sensitivity of acoustic colour can be exploited as well. UXO objects have at least one symmetry axis, whereas natural clutter contacts are generally not symmetric. Because of the amount of detail in multi-aspect acoustic colour, it would be ideally suited to recognize symmetries. To be successful, this would require that a larger angular coverage would be available. This could be obtained by:

- Increasing the angular coverage in a single run. This can only be achieved in the hardware design.
- Combining multiple runs (Figure 5-4). Tracks should be carefully designed such that the object is located in the same range window.
- Circular SAS [12] could be ideally suited as well for increasing the angular coverage. This requires circling around the object in order to obtain information from all angles.

These three ways to improve the angular coverage are shown schematically in Figure 5-5.

Increasing the bandwidth is desirable as well. By increasing the bandwidth, resonances could be observed for a larger range of objects, especially for objects with different sizes.

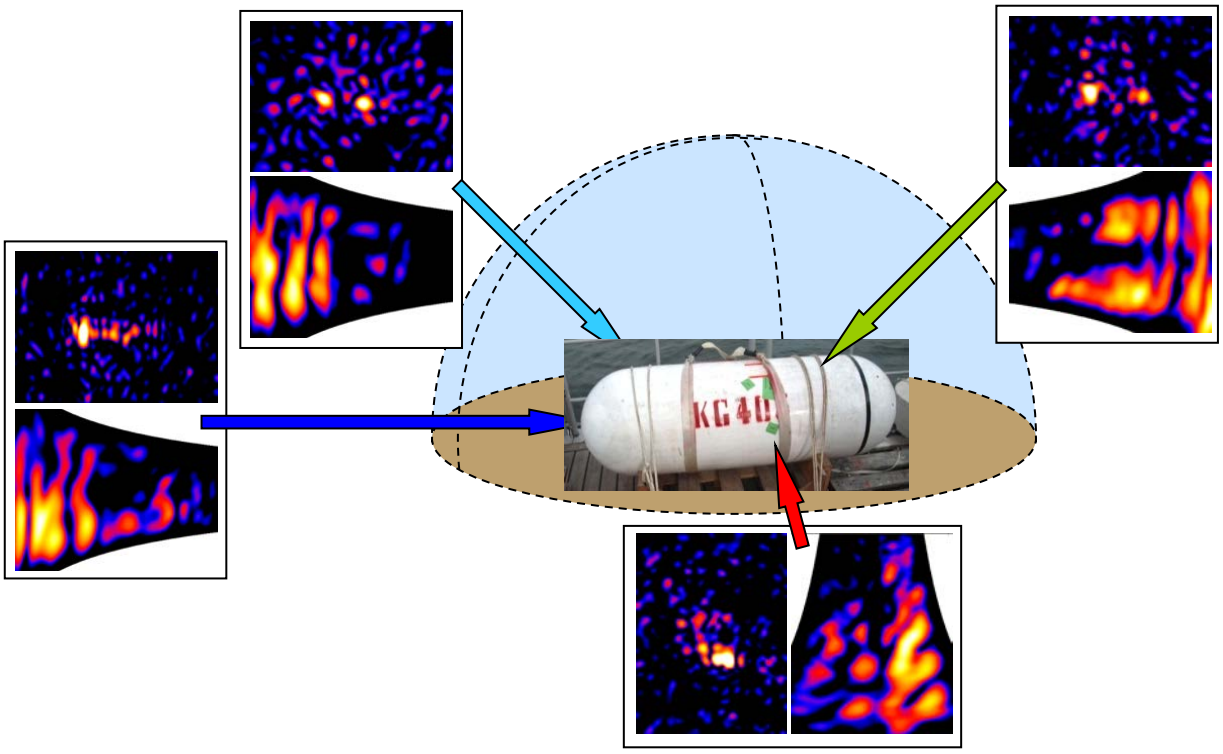

(a)

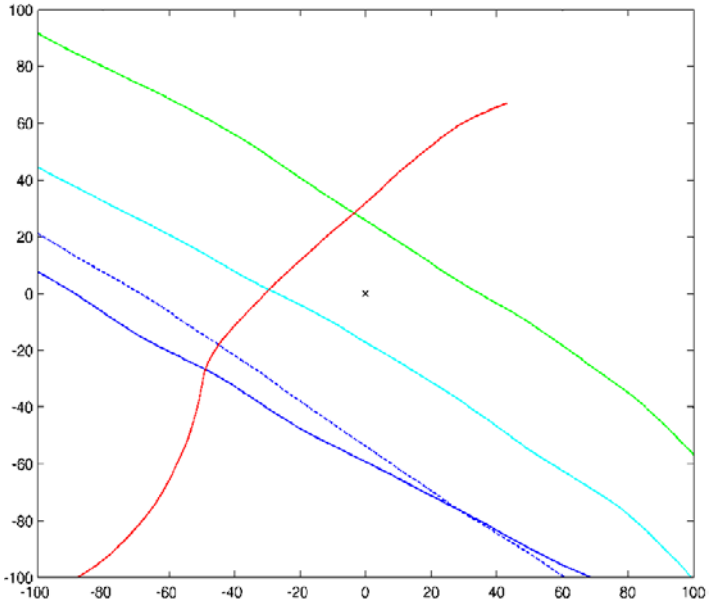

(b)

Figure 5-4 – (a) SAS image snippets and multi-aspect acoustic colour plots from multiple runs of different geometries on the EVA cylinder. The image snippets show an area of 5×4 m around the target, and the acoustic colour plots show dependencies in aspect angle from -45° to 45° relative to the across-track direction and frequencies from 4 to 9 kHz; (b) Track orientations with respect to the target. The colours of the arrows in (a) correspond to the colour of the tracks in (b). The cross indicates the target position.

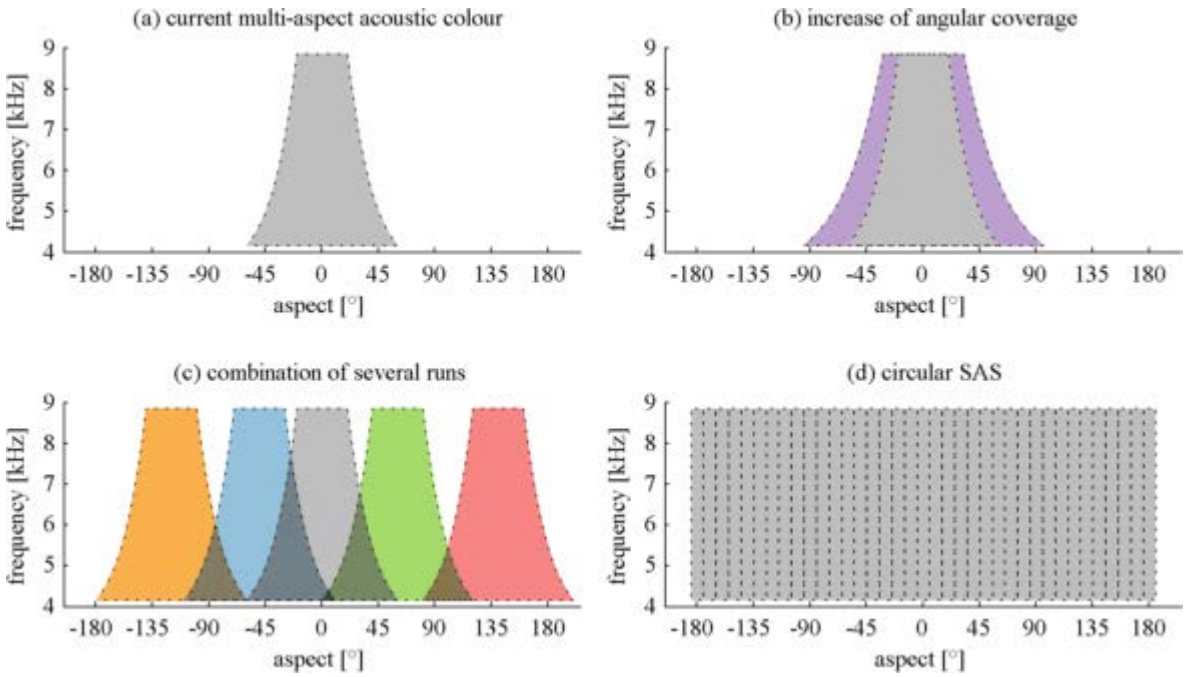

Figure 5-5 – Three ways to increasing the angular coverage of (a) the current multi-aspect acoustic colour plot: (b) by changing the hardware to obtain a larger angular coverage of a single run, (c) by combining several runs into a single plot, and (d) by using circular SAS.

Thus, in order to exploit multi-aspect acoustic colour, we believe that sufficient angular coverage will be the key to success. Then, it should be possible to exploit the angular sensitivity of acoustic colour to discriminate between UXO and natural clutter, especially if the bandwidth of the system is increased as well.

5.4 Clutter Reduction by Contact Classification

The fourth step in the detection chain is the classification of contacts. The primary goal of the classification step is the reduction of clutter contacts while preserving the contacts from relevant UXO. This is, however, a challenging task for several reasons:

- Datasets contain a h igh number of clutter contacts and only a low number of UXO contacts. This means that the contact classification approach has to be capable to deal with imbalanced data both in training and evaluation of the classifier.
- The variability in clutter is very large. Clutter contacts may originate from non-UXO, man-made objects, natural objects and small and large bottom features.
- In the data considered, there is also variability in the deployed UXO objects. Variability is furthermore enhanced because feature values may change with orientation of an object, and could also be influenced by the burial conditions of the object.

Since the functional forms of most of the feature probability distributions are unknown, non-parametric classifiers have been considered for clutter reduction. Two non-parametric classifiers have been tested: a k-Nearest Neighbours (k-NN) classifier and a Parzen classifier. Both methods are kernel density estimators. Since these methods are non-parametric, the challenges as posed above can be handled. Variability in data (both clutter and UXO) can be taken into account by using a s ufficiently large dataset with representative samples. The data imbalance can be compensated, for instance by increasing the weight of the minority class due to specification of an appropriate cost function.

The k-NN classifier is a kernel density estimator, with a uniform, variable sized kernel and a fixed number of samples inside. Given the *k* samples closest to a specific test sample, it is assigned to that class which takes the majority vote [13]. All *k* samples are weighted equally, although the weights assigned to each class can be varied. This weighting makes a k-NN classifier applicable to imbalanced data sets [14]. It corrects for differences in average distance between feature samples

The Parzen classifier is also a kernel density estimator, but with a fixed kernel size and a variable number of samples inside. In general, a Gaussian kernel is used [15][16]. The Gaussian kernel has a single variable parameter for each feature, determining the spread of the kernel, as the standard deviation does in a Gaussian distribution. Additionally, the kernel size can be varied per class. Since the dataset contains much more clutter data, it is assumed that the clutter points in the feature space are closer to each other than the UXO points. In order to correct for this imbalance, the (smoothing) kernel for the UXO should be wider, leading to a more connected kernel density.

An example of Parzen classification is shown in Figure 5-6 for artificial data. Two hardly separable classes are shown, with an imbalanced number of samples. The dataset contains hundred samples from class 1 (circles, representing clutter) and twenty samples from class 2 (crosses, representing UXO). A Parzen classifier with varying kernel sizes has been applied to this training data. In the top row, the kernel sizes of both classes are equally sized. Classification with narrow kernels (a) adapts the feature space to the specific data in the training set, while classification with wide kernels (c) tends to oversmooth the feature space. Adding different weights to the kernels (non-equal kernel sizes for both classes) acts as a penalty function for misclassification, as can be seen in the bottom row. In (d) the class 1 samples (circles) have the widest kernels, leading to confined areas around the class 2 samples (crosses), while this is more or less reversed in (f). As can be seen, a large part of the class 1 samples is correctly classified as class 1, while hardly any of the class 2 sample is misclassified. This illustrates the potential of non-parametric classification for clutter reduction.

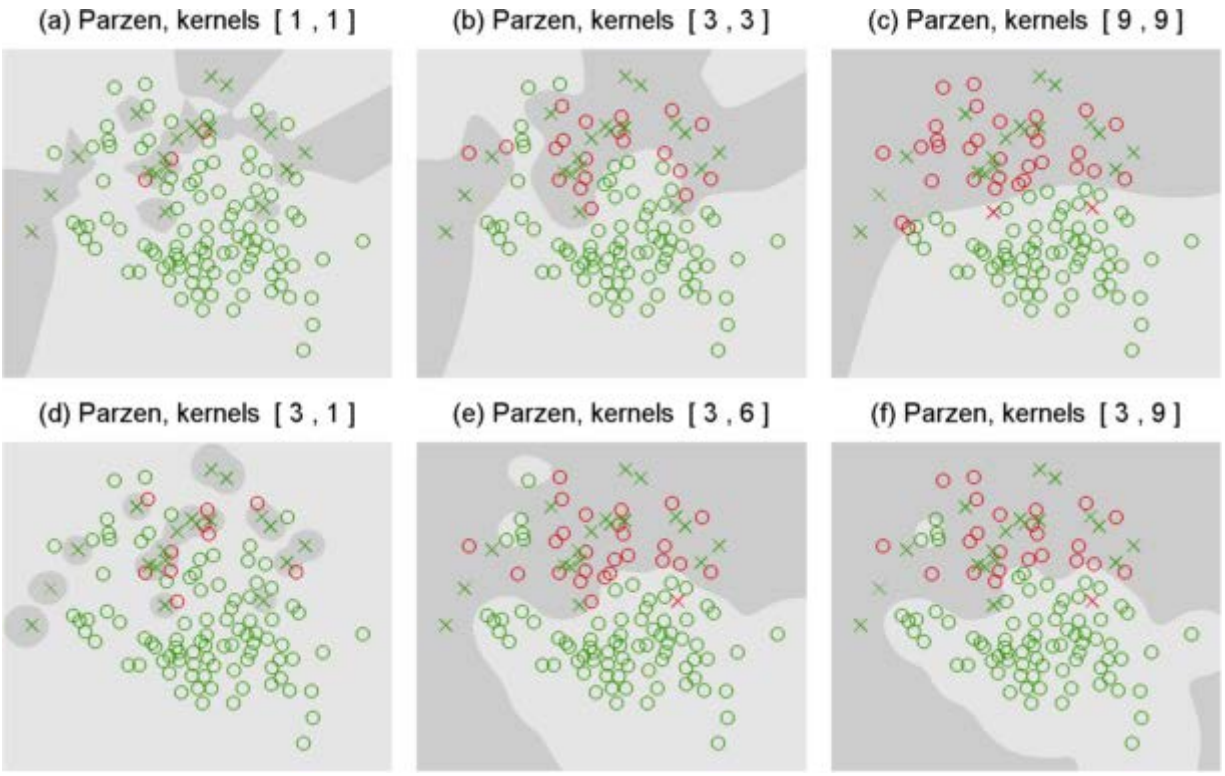

Figure 5-6 – An example of Parzen classification in the training phase, with varying kernel sizes for the two classes. In the title of the subfigures, the kernel sizes are shown between brackets for the two classes respectively. Circles (○) are class 1 samples, while crosses (×) are class 2 samples. Green colour means correct classification, while red means incorrect classification.

6 Results

The developed UXO detection approach has been applied to SAS images from the MUD-2011 trial. In these trials, various targets have been deployed, including UXO targets. In addition to detecting deployed targets, contacts are also obtained corresponding to non-deployed targets. To aid the assessment of the UXO detection, the classification tree shown in Figure 5-1 is further extended including deployed versus non-deployed and man-made versus natural objects, as shown in Figure 6-1. The following categories are used in the analysis:

A. Deployed UXO objects, including shells, bombs, and cylinders. The objective of the detection-classification approach is to distinguish these from all other contacts.

B. Deployed man-made and clutter objects, including a cable, a chain, boulders, and ballast weights. Accurate ground truth information is available on size and shape.

C. Non-deployed clutter objects, mainly natural clutter (ripples). Limited ground truth information is available on these objects.

D. The background, corresponding to parts of the seafloor without any detections, mainly used for reference.

The objective of the UXO detection approach is to separate A from B-D. The non-UXO category is divided into B-D because a classifier is based on the discrimination potential of features between classes, and the nature of the contacts in B and C and the acoustic signatures in D will be completely different. In the following sections, the probability of detection thus refers to the probability of detection of UXO, and false alarms refer to contacts that cannot be associated to UXO objects.

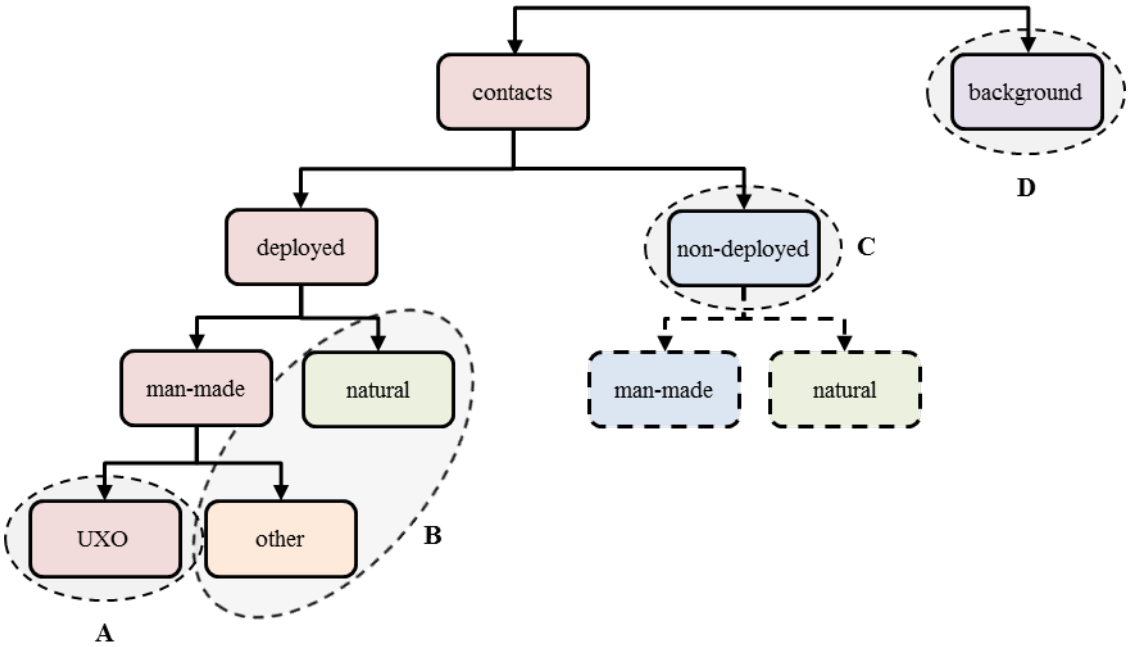

Figure 6-1 – The extended classification tree, with the four categories A-D: UXO, deployed clutter, non-deployed clutter and background, respectively.

6.1 Contact Detection

This section describes the application of the first step of the detection chain, detailed in Section 5.1, to SAS images from the MUD-2011 trial, and investigates how the background removal step influences the contact detection.

The first step in the application of contact detection is to determine tuning parameters of the detector.

The most important parameters are:

- The detection threshold
- The wavelet threshold for the background suppression.

Settings for these parameters have been determined using SAS images acquired in a single run. Consider the SAS image in Figure 6-2. The probability of detection can be obtained by using the information on the locations of the deployed UXO objects. All other contacts are considered to be non-UXO and referred to as 'false alarms'.

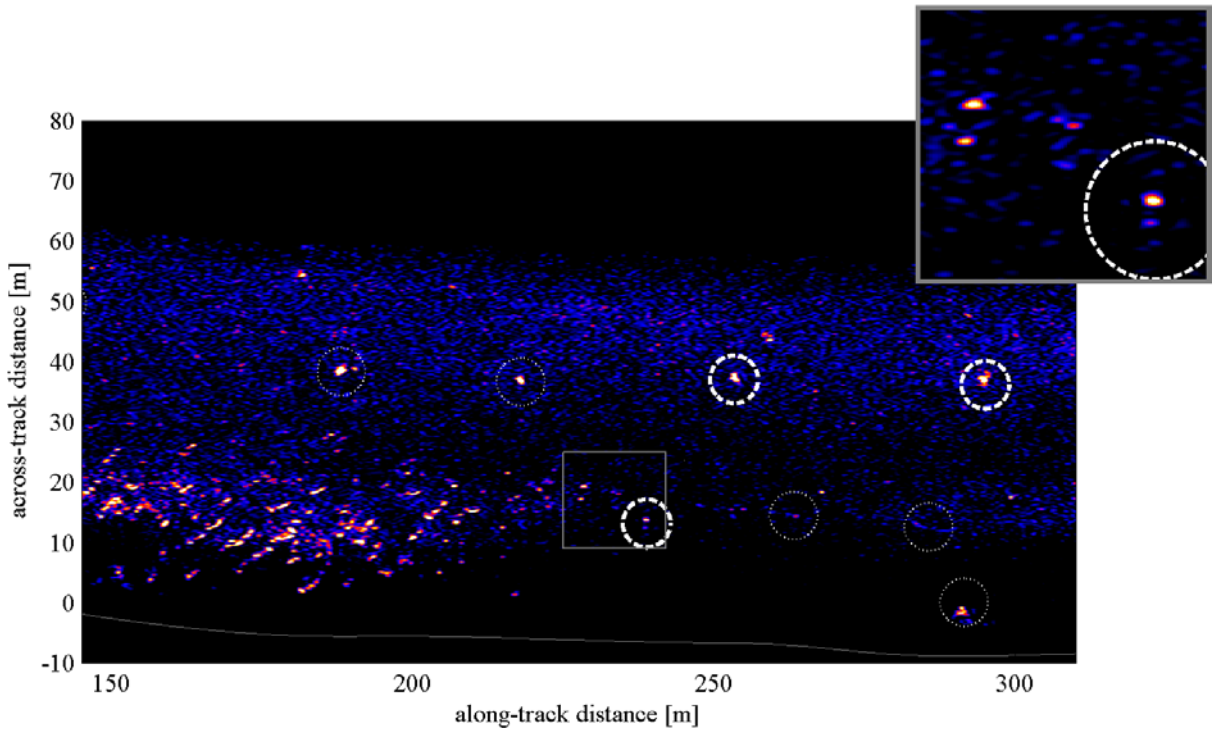

Figure 6-2 – The SAS image used as input for the detection process, where the thick dashed circles denote deployed UXO objects and the thin dotted circles denote deployed non-UXO objects. The inset shows the area marked with a gray square. Background removal has not been applied to this SAS image.

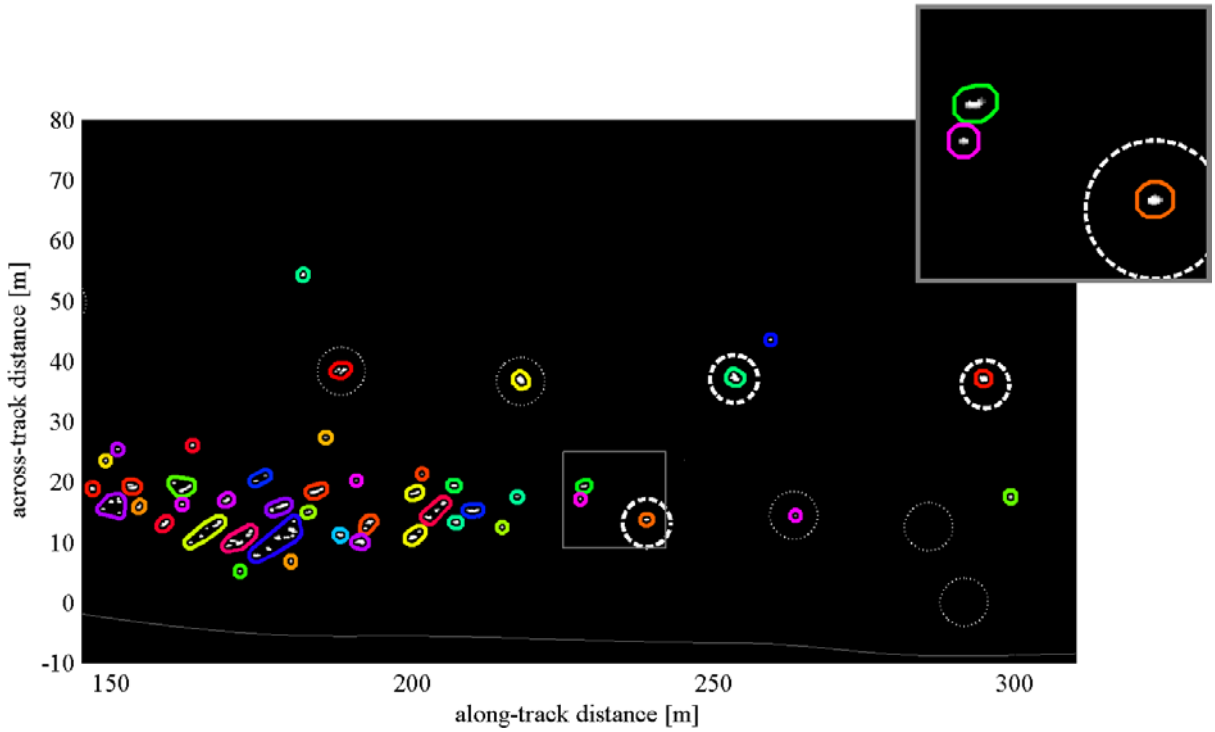

Figure 6-3 – The final result, showing the individual contact clusters, each containing one or more highlights. The inset shows the area marked with a gray square.

Contacts generated by the detector have been analysed in order to tune the detector. Contacts that can be associated to deployed UXO contacts contribute to the probability of detection, whereas all other contacts contribute to the false alarm rate. Figure 6-3 shows an example of the output of the contact detector. It shows that all deployed objects have been detected, and that a finite number of other contacts are generated [17]. The hypothesis is that most of these contacts correspond to the high clutter area in the SAS image caused by seabed features (ripples). Results from the intermediate stages are presented in Appendix B.

Tuning of the detector is further detailed in Figure 6-4. It shows the probability of detection and the false alarm rate for various settings of the normalised detection threshold and normalised wavelet threshold. When background suppression is not applied, the wavelet threshold is 0. In the tuning, other settings are fixed. Specific values for dilation of highlights and clusters have been chosen, such that individual highlights have at least a separation of 1.0 m and highlights separated less than 1.0 m are combined in a highlight cluster. Individual highlight clusters have at least a separation of 4.0 m.

Figure 6-4b shows that the false alarm rate can be reduced by increasing one or both of the thresholds. However, this also leads to a decreased probability of detection (Figure 6-4a). Optimization of the thresholds requires restrictions to the probability of detection and/or the false alarm rate. Here, a high probability of detection is required (e.g. > 90%), because the primary objective of the first step in the detection approach is to maintain all UXO. The lowest false alarm rate is then obtained when the wavelet threshold is chosen to be in the order of 0.5. For further experiments, it is set to 0.45, as in [5]. The detection threshold is kept as a degree of freedom such that ROC curves can be obtained, i.e. to enable the performance comparison with the contact detector for the wavelet threshold equal to 0.

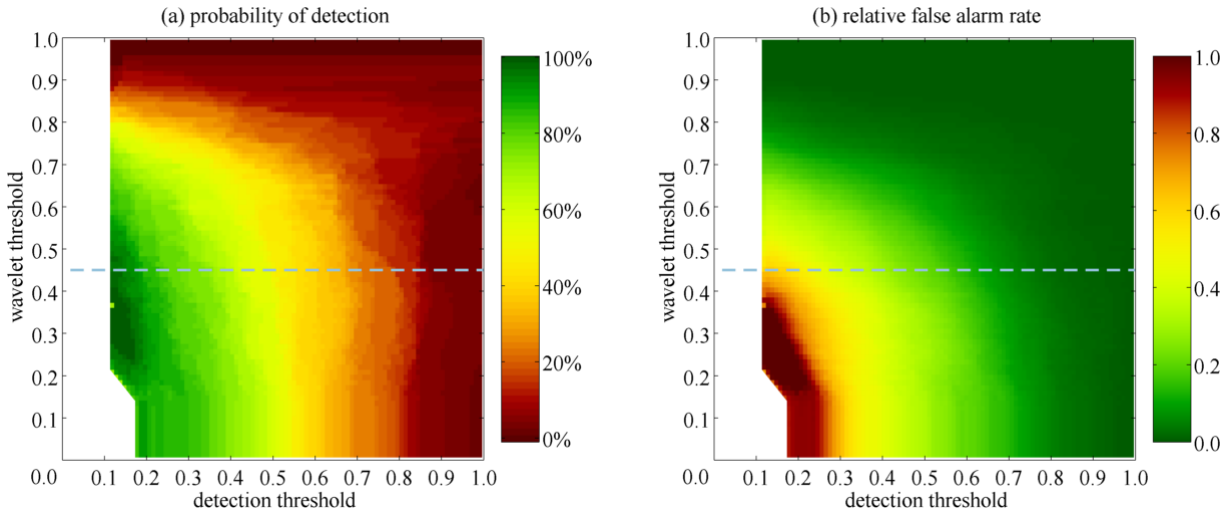

Figure 6-4 – (a) the probability of detection and (b) the relative number of false alarms, both as a function of the wavelet threshold and the detection threshold for the coherent images of 5 runs combined. The dotted line shows the default wavelet threshold value used in [5]. In the blank area, the low detection threshold leads to unreliable results due to fusion of multiple contacts into large areas.

Once the detector has been tuned, the performance of the contact detector can be evaluated by altering the detection threshold. This performance evaluation enables the comparison of the detector applied to SAS images and SAS images with background removed.

To evaluate the performance of the contact detector, Receiver Operating Characteristic (ROC) curves are generated on a more extensive data set, currently containing images of 5 runs from the MUD-2011 data set (Appendix D). These curves show the probability of detection (PD) for UXO versus all other contacts that cannot be associated to UXO contacts, referred to as false alarms (FA). It can be observed that the ROC curve obtained after the background suppression consistently has a better performance than the ROC curve obtained without the application of background suppression. The objective of the second step of the UXO detection approach, the classification step, consisting of signature extraction, feature extraction, and classification, is to further reduce the false alarm rate.

6.2 Extracted signatures

After contact detection, signatures have been extracted and corresponding multi-aspect acoustic colour images have been generated for the different categories given in Figure 5.1:

A. Deployed UXO, 20 contacts from 5 runs.

B. Deployed other objects (man-made clutter and 'natural' clutter objects), 27 contacts obtained in 5 runs.

C. Non-deployed clutter objects that are detected, mostly natural clutter.

D. Background, randomly chosen patches from areas without contacts.

Imagery of the acoustic signatures and corresponding multi-aspect acoustic colour is provided in Appendix A for these four categories. These images provide inspiration for the definition of candidate features (used in Section 6.3).

6.3 Clutter suppression based on image features: First results

The objective of the contact detector is to provide a list of contacts that corresponds to all potential UXO objects. This list can be subsequently investigated in more detail, with the aim to further reduce the clutter. One of the methods to further reduce the clutter is to use a classification approach, i.e. to extract features which have a discrimination potential between UXO and other contacts.

Candidate image features derived from contacts are listed in Section 5.3.1. To initially investigate the discriminative power of features, cumulative empirical distribution functions (EDF) are computed for all considered image-based features for all contacts and for UXO contacts only. Some are shown in Appendix C. When a large difference is observed between the cumulative EDFs for UXO and clutter, the feature has a discrimination potential. However, one cannot draw conclusions on the usefulness of a feature when there is not a significant difference between the EDFs. Then, the *individual* feature cannot be used to separate UXO from natural clutter. One can, however, not exclude that the feature is useful in combination with another feature.

This is further investigated by cross-plotting feature combinations (Figure 6-5). An exploratory analysis of image features reveals that intensity in combination with intensity variation features can be useful. Furthermore, cluster size features combined with intensity variation features could be useful for suppression of clutter contacts as well. In Figure 6-5a, a region can be identified in this with detections that do not correspond to UXO. This feature combination therefore has a potential to reduce the number of false alarms. Another feature combination example is provided in Figure 6-5b. Feature combinations could therefore have a potential for the suppression of natural clutter.

When the red contacts in Figure 6-5 are marked as clutter contacts and consequently removed from the contact list, the remaining number of false alarms will be further reduced. The application of these two rules to the data shown in Figure 6-3, leads to removal of approximately 50% of the contacts. All contacts are shown in Figure 6-6, where the clusters marked as false alarms, are shown in red, while the remaining clusters (both targets and false alarms) are shown in green. Repetition of the whole procedure for varying detection thresholds again leads to a ROC curve, shown in Appendix D. As can be seen, the clutter reduction shows the best results at high false alarm rates.

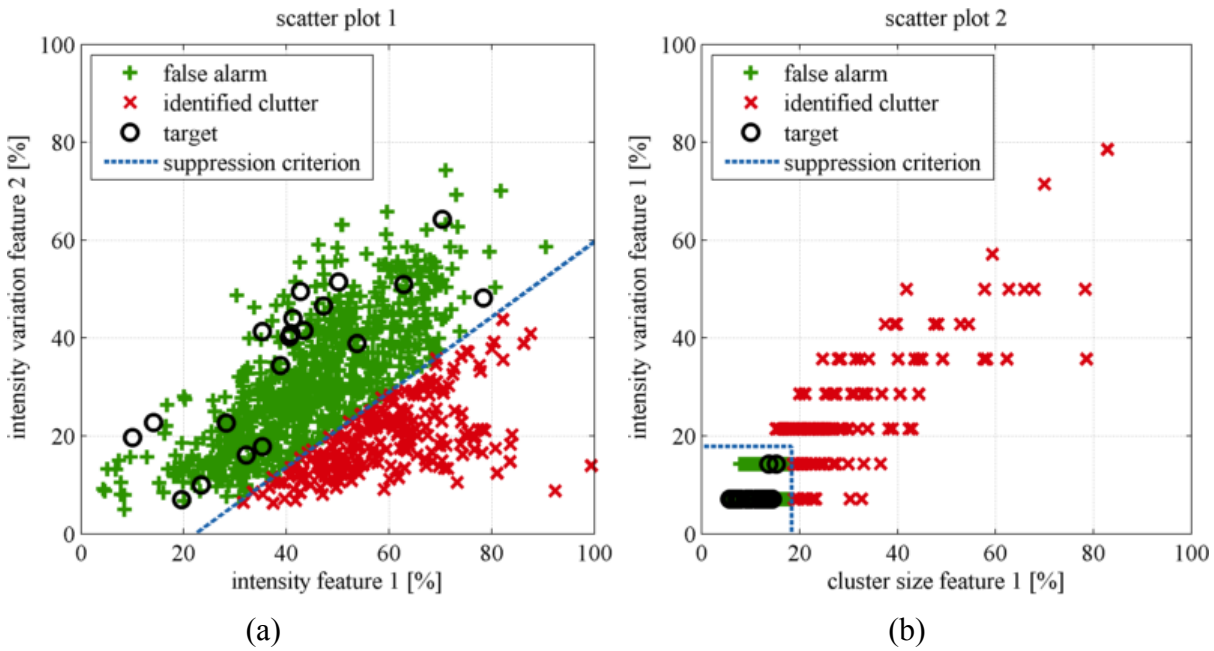

Figure 6-5 – (a) Intensity of all contacts plotted versus intensity variation feature. Contacts below the dashed line can be identified as clutter. (b) Cluster area feature of all contacts versus intensity variation feature. Contacts outside the dashed box can be marked as clutter contacts.

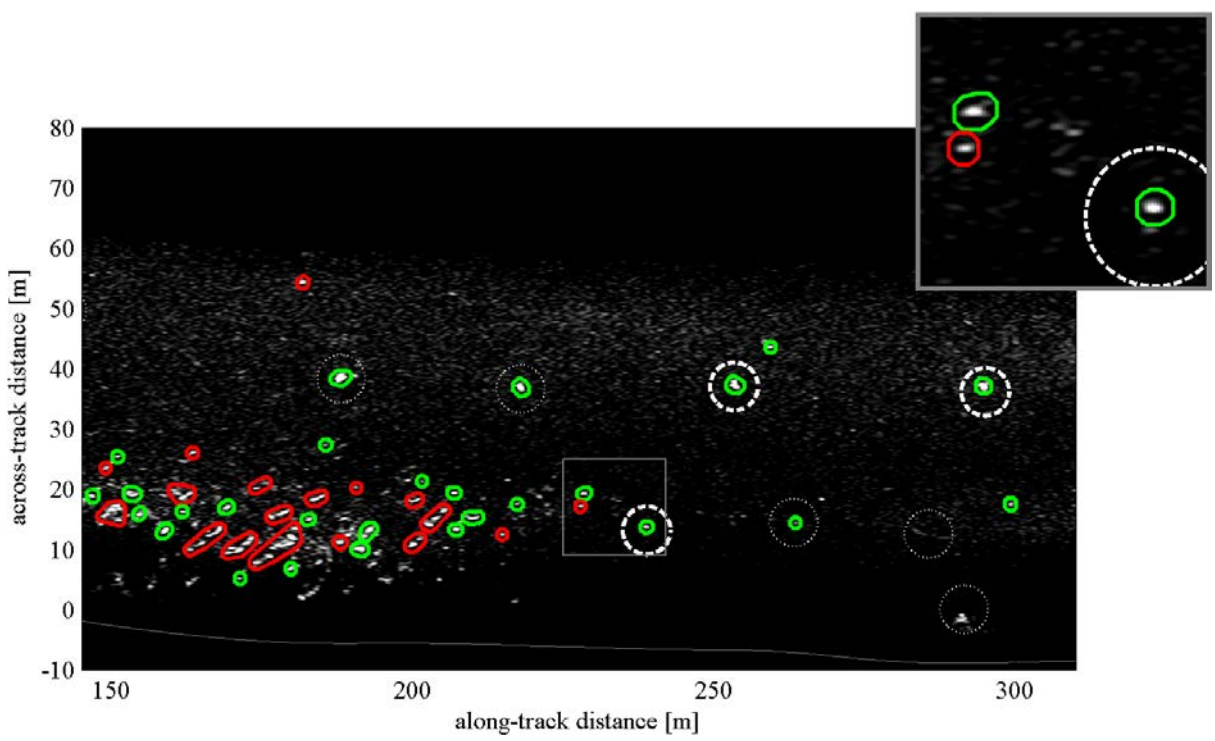

Figure 6-6 – The contacts identified empirically as clutter are marked with a red colour, while the remaining contacts (targets and clutter) are marked with a green colour. The thick dashed circles denote the positions of deployed UXO objects and the thin dotted circles denote the positions of deployed non-UXO objects. The insets show an enlarged part of the images, marked by a gray square.

6.4 Clutter suppression based on image features: Classification results

In the previous section, clutter reduction has been performed empirically. In this section, the proposed k-NN and Parzen classifiers have been applied to the data. As a starting point for classification, the wavelet and detection thresholds are chosen in such a way that none of the UXO contacts are lost. The resulting dataset contains 973 contacts: 953 clutter contacts and 20 UXO contacts. For each of the contacts the aforementioned 6 features are computed: one cluster size feature, one intensity feature, two intensity variation features, and two cluster shape features.

Each class of the dataset was randomly split into training data (67%) and test data (33%). Both classifiers have been applied to the data, and the classification has been repeated 25 times.

A weighted k-NN classifier has been applied for several per-class weightings in order to generate a ROC curve. The weighing provides a penalty for misclassification of UXO with respect to clutter. For the same reason, the Parzen classifier from PRTools [18] has been applied for several combinations of the smoothing kernel and the per-class weighting. An example that illustrates the potential of clutter reduction is shown in Figure 6-7 for the Parzen classifier. For both classification methods, new ROC curves are generated and the mean curves are shown in Appendix D.

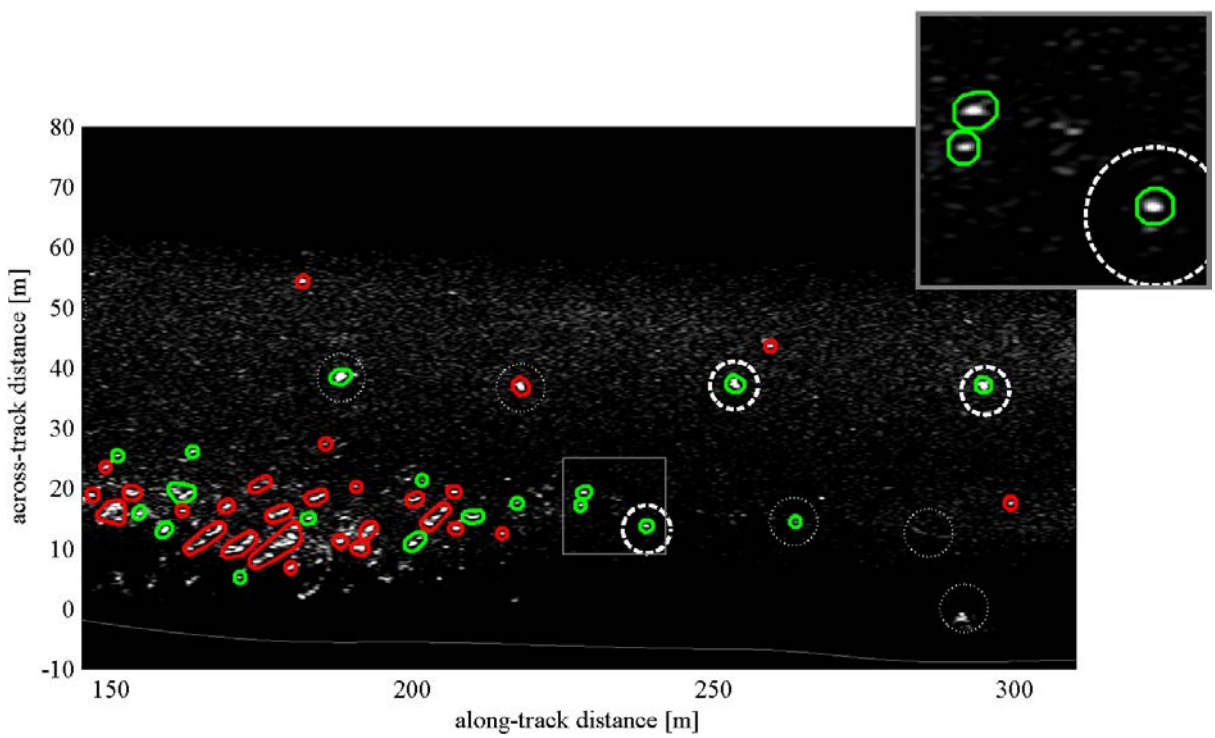

Figure 6-7 – The contacts identified by Parzen classification as clutter are marked with a red colour, while the remaining contacts (targets and clutter) are marked with a green colour. The thick dashed circles denote the positions of deployed UXO objects and the thin dotted circles denote the positions of deployed non-UXO objects. The insets show an enlarged part of the images, marked by a gray square.

7 Discussion

In the previous section, results of a two-step UXO detection approach were presented illustrating the potential of the developed methodology. The following aspects deserve consideration:

- Our objective was to separate UXO objects from all other types of contacts. These are not only contacts corresponding to other deployed man-made objects, but also contacts corresponding to natural clutter.
    - For clutter contacts that do not correspond to deployed objects, there is limited ground truth information available. We decided to include these contacts in the analysis, because these type of contacts will be encountered in practice. It is assumed that these clutter contacts correspond to seabed structures. Nevertheless, on high-frequency REMUS images, we also detected some non-deployed man-made objects in the MUD-2011 trial area.
    - Different types of UXO objects were deployed. This makes the UXO detection and clutter reduction more difficult in comparison to situations in which a single type of UXO objects would be present.
- Only a small amount of UXO objects were deployed. As a consequence, testing and training data for the false alarm reduction step are not fully independent. In the training and evaluation of the false alarm reduction, different views of the same object have been used. Although this is a useful intermediate step to test developed concepts, there is a risk that the false alarm reduction will be over-trained. It is likely that the performance of the false alarm reduction will be lower for objects that are not included in the training data, and thus also for scenarios when it will be applied to new areas. It is, however, not evident how to reduce this risk when data from a single experiment/area are available. One possibility would be to develop augmented reality, in which simulated targets can be included in recorded experimental data. The data provide information on natural clutter contacts, whereas UXO contacts are then imported through simulations [5][20]. The benefit of such an approach has been demonstrated on high-frequency side-scan sonar for the classification of minelike objects [19]. Furthermore, such an approach could also aid to resolve the problem of imbalanced data, i.e. that the number of UXO contacts is limited in comparison to the number of clutter contacts.
- Multi-aspect acoustic colour information has not been used so far. The primary reason is that multi-aspect acoustic colour shows significant variability with object orientation and also with object burial conditions. Because of this variability, it is difficult to develop a feature-based classifier. It should be noted, though, that we believe that acoustic colour information could be relevant if data covering a larger range of aspect angles and frequencies would be available. This statement is supported by the results presented in the SERDP webinar on Acoustic Methods for Underwater Munitions [21].

8 Conclusions

Locating and surveying underwater dumping sites for unexploded ordnance (UXO) is challenging, particularly when they are buried beneath the seafloor sediment. With TNO's MUD low frequency broadband sonar (in the frequency band 4–9 kHz), we have demonstrated that UXO can be detected, even when the UXO are buried in mud. However, besides the contacts that correspond to UXO, there are also a high number of clutter contacts. As a consequence, the use of LF sonar is limited unless a capacity is developed to reduce the number of contacts corresponding to clutter (e.g., natural objects, debris, and seafloor topography). A capability to reduce the number of clutter contacts has been explored in this SERDP MR-2415 project.

As a starting point, we used the data acquisition and (pre)processing chain developed in SERDP MR-2200 project. In this processing chain, the SAS processing takes place, followed by a background suppression step. The next step is detection and false alarm reduction.

A two-step procedure has been followed. First, a contact detector which is capable of detecting objects automatically has been developed. Contacts corresponding to targets and clutter are generated, and these provide input for the second clutter reduction stage. A feature-based approach for clutter reduction has been explored. This two-stage approach has been applied to MUD-2011 data, and the results indicated the potential of this procedure for detecting UXO.

An exploratory investigation on the use of (multi-aspect) acoustic colour has been conducted. Based on this investigation, we concluded that it is difficult to use acoustic colour features for the classification because of limited angle coverage and large feature variability.

9 Way ahead

Limited availability of data with UXO objects is a challenge for the development of an UXO detector. The availability of more data from different representative environments would aid the development of an UXO detection capability. Experimental data not only provide insight in the UXO signatures that can be observed in practice, but also in the acoustic signatures of clutter.

In addition to the availability of more experimental data, it is our opinion that high-fidelity UXO simulations are needed as well. Limitations in the number of UXO objects/contacts in data hinders the development of a robust classifier approach for reducing the number of clutter. On experimental data with limited number of UXO objects, it is generally not feasible to evaluate a classifier on fully independent data, i.e. on objects that are not included in the training set. Commonly, training and evaluation is conducted on different views of the same object. This approach results in a too optimised view on achievable performances. Problems with limitations of number of UXO contacts could be solved by a process to incorporate high-fidelity simulations of UXO objects in the experimental data through augmented reality.

Another potential improvement in false alarm reduction could be achieved by including more or different types of data in this process. By enhancing the aspect coverage and bandwidth, more data becomes available to distinguish UXO from clutter. To achieve this, circular SAS may be

considered, or fusion of data acquired in different runs. One could also consider bi-static acquisition strategies to exploit forward scattering instead of backscattering. The availability of more data should subsequently be used for an updated exploration into useful features for reducing clutter.

Finally, fusion of data acquired with different sensors should be considered. With LF-sidelooking sonar, a relatively high area coverage can be achieved. In combination with other sensors, such as magnetic sensors, regions with high clutter densities could be further explored for the presence of UXO contacts.

10 Publications

A.J. Duijster, A.J. Hunter, R. van Vossen, and A.L.D. Beckers, *Towards Automatic Target Recognition in Low-frequency Sub-sediment Sonar Imagery*, UAC, Rhodos, Greece, June 2014. [Appendix D]

References

[1] S.G. Schock, A. Tellier, J. Wulf, J. Sara, and M. Ericksen, *Buried Object Scanning Sonar*, IEEE Journal of Oceanic Engineering, 26, p. 677-689, 2001.

[2] A.L.D. Beckers, R. van Vossen, S.P. Beerens, B.A.J. Quesson, M. Zampolli, J.C. Sabel, and J. Janmaat, *MUD Trial Final Report*, TNO-DV 2010 A004, 2010.

[3] A.L.D. Beckers, B.A.J. Quesson, and R. van Vossen, *MUD Trial Phase 2 – Final Report*, TNO-DV 2011 A343, 2011.

[4] J.A. Bucaro, B.H. Houston, M. Saniga, L.R. Dragonette, T. Yoder, S. Dey, L. Kraus, and L. Carin, *Broadband Acoustic Scattering Measurements of Underwater Unexploded Ordnance (UXO)*, Journal of the Acoustical Society of America, 123, p. 738-746, 2008.

[5] K.L. Williams, S.G. Kargl, E.I. Thorsos, D.S. Burnett, J.L. Lopes, M. Zampolli, and P.L. Marston, *Acoustic Scattering from a Solid Aluminum Cylinder in Contact with a Sand Sediment: Measurements, Modelling, and Interpretation*, Journal of the Acoustical Society of America 127, p. 3356-3371, 2010.

[6] A. Tesei, M. Zampolli, and G. Canepa, *At-sea Measurements of Acoustic Elastic Scattering by a 1.5 m Long Cylinder Made of Composite Materials*, 2nd International Conference on Underwater Acoustic Measurements, Crete, Greece, 2007.

[7] A. Bellettini, and M.A. Pinto, *Theoretical Accuracy of Synthetic Aperture Sonar Micronavigation using a Displaced Phase-centre Antenna*, IEEE Journal of Oceanic Engineering 27, p. 780-789, 2002.

[8] A.J. Hunter, and R. van Vossen, *Acoustic Signatures of Underwater UXO Measured by Low Frequency Broadband SAS*, UAC, Corfu, Greece, June 2013.

[9] A.J. Hunter, R. van Vossen, B.A.J. Quesson, and A.L.D. Beckers, *Detection of Underwater UXOs in Mud*, SERDP Project MR-2200, April 2013.

[10] A.J. Hunter, and R. van Vossen, *Sonar Target Enhancement by Shrinkage of Incoherent Wavelet Coefficients*, Journal of the Acoustical Society of America 135(1), January 2014.

[11] T.M. Marston, P.L. Marston, and K.L. Williams, *Scattering Resonances, Filtering with Reversible SAS Processing, and Applications of Quantitative Ray Theory*, OCEANS 2010, p. 1-9, 2010.

[12] T.M. Marston, J.L. Kennedy, and P.L. Marston, *Coherent and Semi-coherent Processing of Limited-aperture Circular Synthetic Aperture (CSAS) data*, OCEANS 2011, p. 1-6, 2011.

[13] R.O. Duda, P.E. Hart, and D.G. Stork, *Pattern Classification*, 2nd Edition, Wiley, 2000.

[14] H. Dubey, and V. Pudi, *Class Based Weighted K-Nearest Neighbor over Imbalance Dataset*, Advances in Knowledge Discovery and Data Mining, Lecture Notes in Computer Science 7819, pp 305-316, 2013.

[15] R.P.W. Duin, *On the Choice of Smoothing Parameters for Parzen Estimators of Probability Density Functions*, IEEE Transactions on Computers 25, p. 1 175-1179, 1973.

[16] M.A. Kraaijveld, *A Parzen Classifier with an Improved Robustness against Deviations between Training and Test Data*, Pattern Recognition Letters 17, p. 679-689, 1996.

[17] A.J. Duijster, A.J. Hunter, R. van Vossen, and A.L.D. Beckers, *Towards Automatic Target Recognition in Low-frequency Sub-sediment Sonar Imagery*, UAC, Rhodos, Greece, June 2014.

[18] R.P.W. Duin, P. Juszczak, P. Paclik, E. Pekalska, D. de Ridder, D.M.J. Tax, and S. Verzakov, *PRTools4.1, A Matlab Toolbox for Pattern Recognition*, Delft University of Technology, 2007.

[19] C. Barngrover, R. Kastner, S. Belongie, *Semisynthetic versus real-world sonar training data for the classification of mine-like objects*, IEEE Journal of Oceanic Engineering 40, 48-56, 2015.

[20] D.S. Plotnick, P.L. Marston, K.L. Williams, A.L. Espana, *High-frequency backscattering by a solid cylinder with axis tilted relative to a nearby horizontal surface*, Journal of the Acoustical Society of America 137, 470-480, 2015.

[21] J.L. Bucaro, K.L. Williams, *SERDP Webinar on acoustic methods for underwater munitions*, February, 2015.

SERDP Munitions Response
MR-2415 Final Report Addendum
June 2015
Principal Investigator – Dr. Robbert van Vossen
TNO

The above mentioned Final Report was produced as a requirement of the Strategic Environmental Research and Development Program (SERDP), Project #2415, UXO Detector for Underwater Surveys Using Low-Frequency Sonar.

This page contains information regarding Appendices A through D of the SERDP Munitions Response MR-2415 Final Report detailing work conducted by Dr. Robbert van Vossen of TNO. Appendices A through D are categorized under Distribution Statement D and are only available to Department of Defense and U.S. DoD contractors. Contents of each appendix are detailed below:

Appendix	Content
Appendix A	Acoustic Signatures
Appendix B	Intermediate Steps of the Detection Chain
Appendix C	Cumulative Empirical Distribution Plots of Features
Appendix D	ROC curves

If interested in requesting copies of these appendices, please contact the SERDP Munitions Response Program Manager via email:

Dr. Herb Nelson
Munitions Response Program Manager
Strategic Environmental Research and Development Program (SERDP)
herb.nelson@nrl.navy.mil

Appendix E – UA paper

The following paper [17] was presented at the Underwater Acoustics Conference 2014.

TOWARDS AUTOMATIC TARGET RECOGNITION IN LOW-FREQUENCY SUB-SEDIMENT SONAR IMAGERY

Arno Duijster[a] , Alan Hunter[a], Robbert van Vossen[a], and Guus Beckers[a]

[a]TNO (Netherlands Organisation for Applied Scientific Research

PO Box 96864, 2509 JG The Hague, The Netherlands, +31 88 866 0902, alan.hunter@tno.nl

Abstract: Detection of unexploded ordnance is challenging in the underwater environment, particularly when object burial occurs. A capability to detect buried targets has been demonstrated previously using TNO's MUD low frequency sediment-penetrating sonar and other similar sonars. However, the high clutter rates encountered in practice have the potential to impose severe operational limitations in absence of a robust capability to distinguish targets from clutter. To this end, we are taking the initial steps towards development of an automatic target recognition algorithm for detecting targets and suppressing clutter in low-frequency sub-sediment sonar imagery. The initial implementation presented in this paper uses a previously developed wavelet shrinkage algorithm to suppress the background reverberation, followed by automatic thresholding and segmentation to isolate individual seafloor objects for subsequent extraction of their acoustic signatures. We show preliminary detection results from the MUD-2011 data set.

Keywords: *Synthetic aperture sonar, buried targets, automatic target recognition*

1. INTRODUCTION

In an underwater environment, the detection of unexploded ordnance (UXO) is difficult when the objects are partially or fully buried. Sine burial frequently occurs, either in conditions with soft seafloors (silt or mud) or due to sediment transport, methods need to be developed to aid the detection of UXO in such conditions. Broadband low-frequency sonars are a promising technology for sub-sediment imaging and detection of buried objects [1],[2],[3]. The low frequencies enable penetration into the seafloor sediment, while broadband signals facilitate potential classification of objects based on the multi-aspect acoustic colour. TNO's hull-mounted side-looking low-frequency synthetic aperture sonar (referred to as the MUD sonar) has been developed for this purpose [1].

UXO detection is challenging with the MUD sonar due to high levels of reverberation, the presence of clutter, and low target echo amplitudes. To develop an effective target detector, a two-stage approach is proposed:

1. A contact detector for selecting objects of interest, including UXO *and* clutter.
2. A classifier which extracts features from the contacts with the objective to distinguish UXO from the clutter.

This paper focuses on the first stage detector mentioned above. For this purpose, dedicated image processing has been developed and applied to experimental data. The image processing is summarised in Section 2, and consists of synthetic aperture processing, followed by target/background separation. Section 3 discusses the contact detector. The results on data acquired by the MUD system are discussed in Section 4 and conclusions are drawn in Section 5.

2. MUD SONAR AND PROCESSING CHAIN

The MUD sub-sediment imaging sonar is a side-looking system with flexible tilt angle. It has a 16-element vertical array, enabling the signal to reverberation ratio to be improved by the suppression of multipath reverberation. It has an accurate navigation system comprised of RTK-GPS and INS, and a horizontal array to aid the synthetic aperture processing. The system is capable of operating in the frequency range between 1 kHz and 30 kHz [1].

2.1. Multipath Suppression and SAS Processing

In shallow water, strong multipath interference can mask the echoes from targets and corrupts their acoustic signatures, particularly at longer ranges. This adversely affects detection and classification performance. Vertical beamsteering is applied for the mitigation of the multipath interference as described in [4].

Synthetic aperture sonar (SAS) uses coherent processing of the echo data to attain high resolution in the along-track direction. This is especially important for low frequency sonars, which typically have poor resolution due to their wide beams. The applied time domain back-projection SAS processing includes motion compensation, and is described in detail in [4].

2.2. Target / Background Separation by Incoherent Wavelet Shrinkage

Effective detection of targets in background reverberation noise requires a robust method, which is able to discern even weakly scattering objects in a highly reverberant seafloor. A coherence metric can be used to separate the targets from the background. In [5], a coherence metric was derived to determine the similarity of wavelet coefficients between independent looks, i.e. different images of the same scene with statistically independent noise realizations. It is assumed that a high coherence corresponds with the reverberation-free measurements of targets, while the low coherence contributions are assumed to correspond to reverberation (background). By weighting the image according to this coherence, a separation can be made between targets and background. A thorough description of this method can be found in [5].

3. TARGET AND CLUTTER DETECTOR

The detection of buried targets in sonar images is difficult due to the presence of clutter and reverberation. This is illustrated in Figure E-8a, where the ground truth positions of deployed objects are overlaid on a SAS image from the MUD sonar. Here, we present an automated approach to obtain contacts corresponding to both targets of interest and clutter. This is a first basic step towards automatic target detection. Once these contacts are obtained, more detailed information can be derived using dedicated processing, for example based on features of the multi-aspect acoustic colour [4].

The operation of the contact detector is described in Figure E-9 and consists of the following steps:

1. Apply a fixed threshold to the SAS image leading to a binary image.
2. Apply morphological opening and closing operations to obtain pixel clusters.
3. Generate a list of contacts from the pixel clusters.

To assess the performance of this basic detector, the known ground truth positions of targets are associated with the detected contact positions, taking into account an assumed maximum distance between a contact and the ground-truth position. The purpose of this step is to estimate the probability of detection (i.e. whether contacts are generated corresponding to deployed targets) and the number of false alarms (i.e., the remaining non-associated contacts). Since the number of false alarms in an image is a non-normalised number, depending on the size of the image, it can be normalised by the area. Here, the normalisation area is chosen to be a square swath-width (50 m × 50 m). A receiver operating characteristic (ROC) curve is obtained by

plotting the proportion of detected targets (true positive rate) versus the normalised number of false alarms (false positive rate).

After sonar data processing, a SAS image is obtained as illustrated in Figure E-8a. This image is naturally used as input for the detection process (shown in the flow diagram in Figure E-9 as a dashed line) to provide a baseline performance estimate. However, because the image contains a clearly visible background and clutter contacts, it is anticipated that the thresholding will not be effective and will result in many false alarms. Therefore, a s econd pre-processing step is considered, whereby a coherence filtering operation (incoherent wavelet shrinkage) is used to separate the image into coherent and incoherent components; these are assumed to correspond with the objects and background reverberation, respectively. The coherent part is shown in Figure E-8b.

In the next section, the performance of the detector is evaluated when it is directly applied to the SAS image, and when the incoherent wavelet shrinkage technique is included in the processing chain to remove the background reverberation.

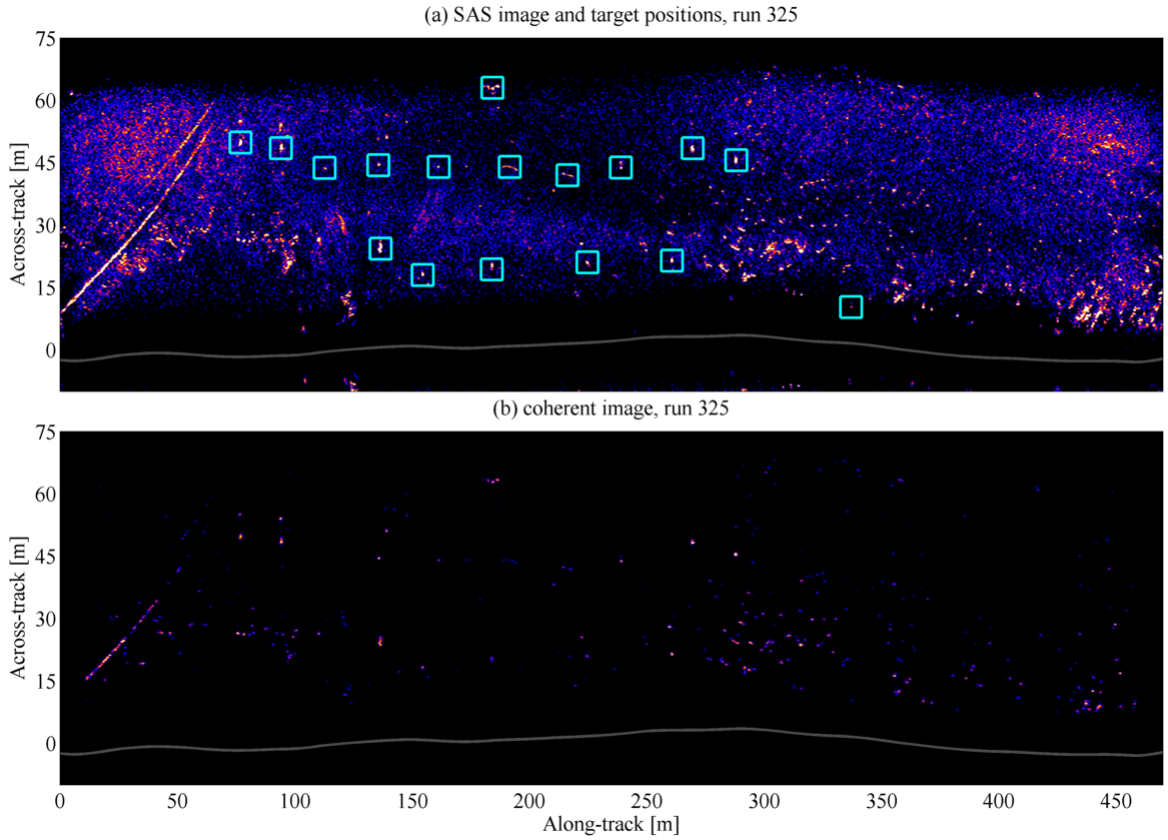

Figure E-8 (a) SAS image and ground truth target positions, marked with a square; (b) the same image, after wavelet shrinkage, showing the coherent part.

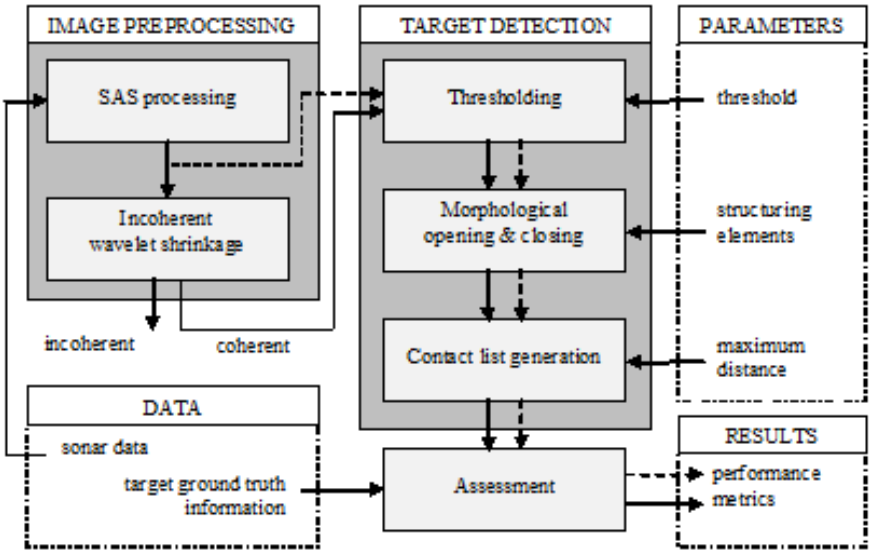

Figure E-9 Flow diagram of the image processing and target detection chain.

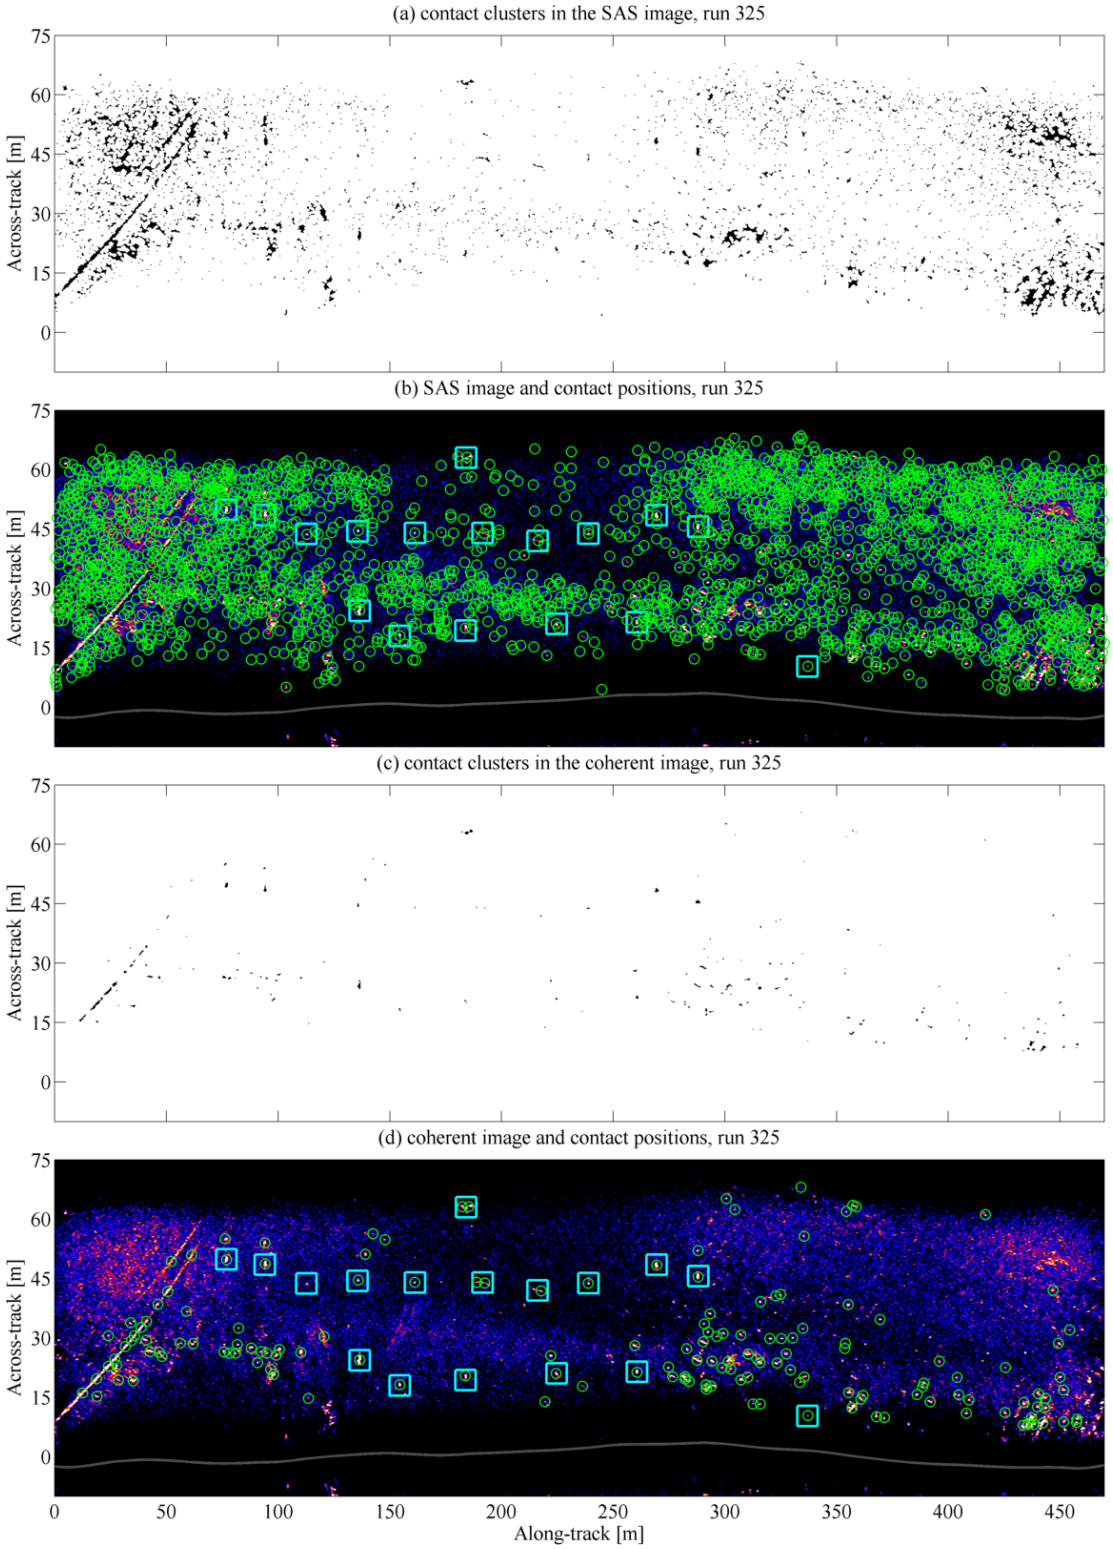

Figure E-10 The binary contact clusters (after the morphological opening and closing step) for both the SAS image (a) and the contacts image (c), and the resulting detections overlayed on the SAS image (b) and the coherent image (d). Ground truth positions are marked with a cyan square, while contact detections are marked with a green circle. The colour scale is identical in the images and is normalised between 0 and 1.

4. RESULTS

The performance of the detection chain is influenced by the choice of a few parameters (shown on the right-hand side of Figure E-9). Here, the structuring elements are fixed with a circular element of 3 pixels in diameter for the opening and a circular element with 7 pixels in diameter for the closing. The maximum distance for the association between a ground truth target position and a detected contact is chosen to be 2 m. The only remaining parameter is the detection threshold, which affects the trade-off between detection performance and false alarm rate. For normalised images, this is varied between 0 (all pixels) to 1 (no pixels). Default parameters are used in the wavelet shrinkage algorithm [5].

Example results using a detection threshold of 0.05 are shown in Figure E-10. The top image (a) shows the thresholded SAS image after the morphological operations. As can be seen, clutter from the background is still present, and these contacts contribute to the false alarms. In (b) the resulting detections are overlaid on top of the SAS image. The large number of false alarms confirms the prediction that a second pre-processing step is necessary. The same procedure is repeated for the coherent image, where the background clutter has been suppressed. The resulting binary image (c) and the overlay on top of the SAS image (d) show a much better result. The example reveals that the automated target detection chain applied directly to the SAS image has a 100% probability of target detection, but at the cost of 276 false alarms per square swath width (i.e. 2500 m^2). When applied to the coherent image, the probability of detection is 94% (16 out of 17), but with a much lower false alarm rate of 15 false alarms per square swath width.

The procedure can be repeated for a range of detection thresholds to yield a ROC curve. The ROC curves are shown in Figure E-11 for two separate runs, each containing the same 17 targets, with and without the background removal. As can be seen, the number of false alarms is significantly reduced using the proposed method, although at the cost of some mis-detections.

5. CONCLUSIONS AND DISCUSSION

We have described and demonstrated a basic contact detector for a low-frequency sediment-penetrating sonar. The detector is sensitive to UXO and clutter and will provide the first stage in a future automatic detector. The performance of the proposed detector was demonstrated with ROC curves and compared with the performance of a basic threshold detector.

We have shown that a substantial improvement in false alarm suppression is achieved by applying a wavelet shrinkage pre-processing step before applying the detection threshold. Although the false alarm suppression is improved by this approach at lower sensitivities (roughly 1 order-of-magnitude better), the probability of detection is affected at higher sensitivities, attaining only 90-95% detection probability at the highest sensitivities compared to 100% for the simple detector (at the cost of a very high false alarm rate). This suggests several improvements that will be investigated in future work:

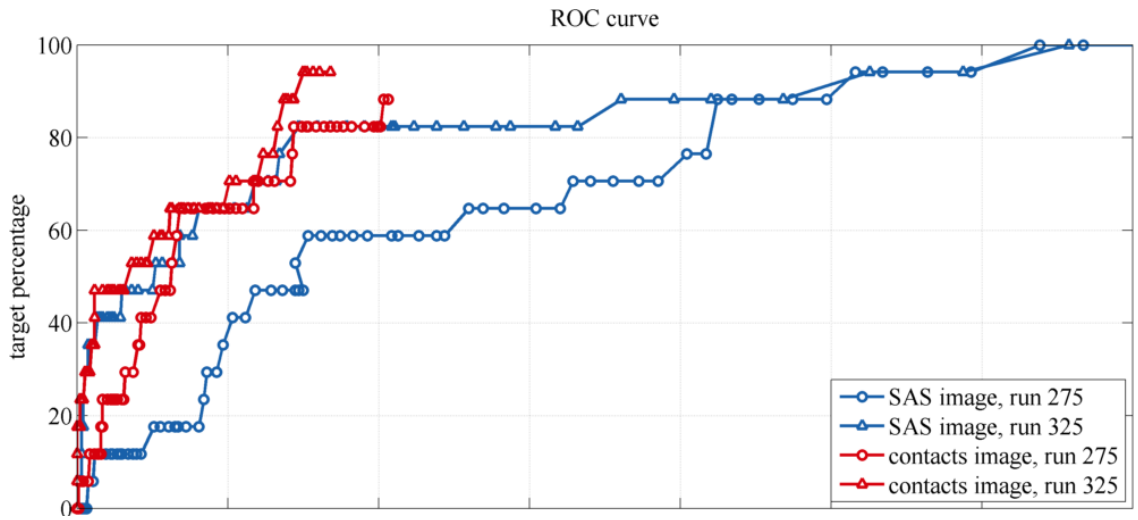

Figure E-11 The ROC curves for the SAS and post-processed images for two separate runs.

1. Optimisation of the detector parameters, including settings of the wavelet shrinkage algorithm, morphological operators, and clustering; and 2. using a combination of results from the processed and non-processed images. Moreover, using the contact detector, we intend to extract signatures from many UXO and clutter contacts to establish a robust feature set for UXO detection.

6. ACKNOWLEDGEMENTS

This work was funded by the US Dept. of Defense Strategic Environmental Research and Development Program. The authors thank the Royal Netherlands Navy for funding and operational support for the MUD trials.

REFERENCES

A.L.D. Beckers, R. van Vossen, and G. Vlaming, Low-frequency synthetic aperture sonar for detecting explosives in harbours, *Sea Technol.* 53(1), March 2012.

P.L. Nielsen, R.D. Hollett, G. Canepa, and W.L.J. Fox, *Unique low-frequency mine hunting and seabed characterization sonar*, UAC, Corfu, Greece, June 2013

S.G. Schock, J. Wulf, G. Quentin, and J. Sara, *Synthetic aperture processing of buried object scanning sonar data*, IEEE Oceans, Washington DC, Sep 2005

A.J. Hunter and R. van Vossen, Acoustic signatures of underwater UXO measured by low frequency broadband SAS, *UAC*, Corfu, Greece, June 2013.

A.J. Hunter and R. van Vossen, Sonar target enhancement by shrinkage of incoherent wavelet coefficients, *J. Acoust. Soc.Am.* 135(1), January 2014.