AUTOMATIC CHANGE DETECTION FOR THE MONITORING OF CLUTTERED UNDERWATER AREAS

Enrique Coiras[a], Johannes Groen[a], David Williams[a], Benjamin Evans[a], Marc Pinto[a]

[a] NATO Undersea Research Centre (NURC).

Enrique Coiras: Coiras@nurc.nato.int, +39-0187527417 (tel), +39-0187527330 (fax).
Postal address: NATO C/O NURC. Viale San Bartolomeo 400. 19126 La Spezia, Italy.

Abstract: *Automating the surveillance of cluttered underwater areas (such as ports) poses several problems, the main one being the high number of contacts that any automatic target recognition (ATR) system will produce when surveying the area. In this paper an approach based on Change Detection which is capable of coping with the expected high number of potential threats is proposed for the surveillance of cluttered areas. To that end an automatic detector is coupled with a Data Association mechanism in order to determine the location of previously undetected contacts which may constitute new threats. The proposed system is tested on a real scenario using synthetic aperture sonar (SAS) data, and the influence of conducting surveys in different orientations on the change detection results is studied.*

Keywords: *change detection, data association, Automatic Target Recognition, ATR, port protection, IED, UWIED.*

1. INTRODUCTION

Over recent years much attention has been given to terrorist threats to Ports and Harbours which is not surprising given the potential economic impact of closing a major port. As an example, the West Coast Labor slowdown in the fall of 2002 cost an estimated 1.95B$/day [1][2] and it remains the case that the majority of world trade is still carried by sea, with, for example, 90% of both US and European imports and exports still transit by sea [3][4]. Many different threats [5] have been envisaged to ports and harbours, including mines, IEDs, suicide boats, underwater swimmers, and exploding fuel tankers to name but a few. This paper concentrates on technology to counter the underwater mine and IED threat. According to [6] 407 million metric tons of goods are transported through the port of Rotterdam each year.

Mines and underwater IEDs are a true asymmetric threat [1]. They are cheap to produce, typically less than $30k for even the most advanced modern systems, whilst basic underwater IEDs can be manufactured simply from a very low technology base. Conversely, countering this threat can require far more sophisticated and capable platforms. Finding mines, particularly in areas where the seabed is highly textured or cluttered with either natural or man made objects can be a difficult and time consuming task. Obviously the faster a closed port can be re-opened the less the economic impact, and thus technology which can be used to quickly and robustly ensure that an area is free of mines is of great benefit.

2. APPROACH TO CHANGE DETECTION

One potential approach to achieving this is to perform regular sonar surveys of the seabed and to then compare consecutive surveys to determine whether / where the seabed has changed—and thus where any new targets have been laid. The change detection analysis described here is based on the SAS images that were collected by the MUSCLE AUV before and after deployment of a control set of known targets. We will refer to those times as Day1 and Day2. In order to damage a ship mines or IEDs need to contain a certain amount of explosive. Potential threats can therefore be determined by focusing on previously unseen objects of a given minimum size. This requires, first, the detection of objects of at least that size—a task that can be performed automatically by a target detector [7] or ATR system [8]—and, second, a way to determine what detected contacts haven't been observed before. For this latter task, a contact matching method that compares a list of potential contacts to a reference list of existing contacts is required. In this paper we use NURC's previously developed contact matching algorithm [9][10] for this purpose. This method was used to good effect to match contacts from adjacent sonar legs in order to remove residual navigation errors and enable new contacts to be identified. The ability of the algorithm to perform satisfactorily on a cluttered seabed, especially for data sets where surveys are performed using different survey directions has however never been tested.

This paper addresses these aspects and also examines the performance of the algorithms when applied to high resolution SAS data gathered using NURC's MUSCLE AUV during the COLOSSUS 2 sea trial, performed jointly between NURC and the Latvian Navy in April / May 2008.

3. DETECTION BASED ON A SLIDING TEMPLATE

The detection algorithm is based on correlation using a template representing an idealized target response [7]. The template is divided into three areas: the background (value 0); a rectangle representing a highlight (value 1) and a rectangle representing shadow (value -1). Dimensions of the highlight and shadow areas can be set according to object dimensions and its relative position to the sonar. For the results presented in this paper the size of the highlight rectangle was fixed to 1m along track and 0.5m across track. The along-track size of the shadow's rectangle was set to 0.25m, while its length is dynamically calculated from the imaging geometry assuming a sonar altitude of 13m and a target height of 0.3m. An example of the results obtained by this detector is presented in Figure 1, where its response to the presence of two cylindrical objects is shown. Locations of likely targets are returned by the detector, along with a small image snippet of the area. These snippets can later be reviewed by an automatic classification system or a human operator in order to discard false alarms or take immediate action to obvious threats (known mine models, for instance).

The performance of the selected detector was measured against a set of 80 ground-truthed image snippets from a previous SAS trial (in Framura, Italy), of which 15 snippets contained targets. The output of the detector is a score for every image pixel, related to the likelihood of the pixel belonging to a target. Cutting off the detector's output for different thresholds produces different detection performances, as shown in the Receiver Operator Characteristic (ROC) curve of Figure 1, right.

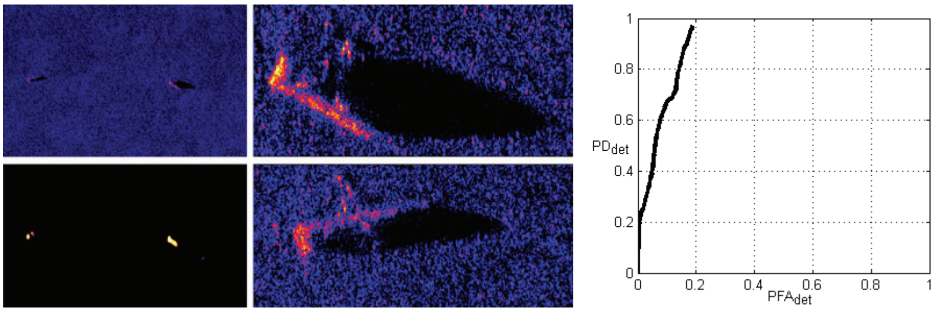

Fig.1: Left: the detector applied to a SAS image containing two cylindrical targets, and the image snippets corresponding to the locations of likely targets. Right: ROC curve for the detector when varying the score threshold.

For the work described in this paper, the threshold was selected to maximize the probability of detection. The threshold that takes us closer to perfect detection using the scores from the Framura study (with a probability of detection PD_{det}=0.97) results in a probability of false alarm PFA_{det}=0.19. Note that these figures are just approximations, but nevertheless give an idea of the behaviour that can be expected from the detector. We therefore used these performance scores as nominal values for the application of the detector to the COLOSSUS2 SAS mission data.

4. CONTACT MATCHING BY RIGID DATA ASSOCIATION

Application of the detector described in the previous section to the sonar images of a given mission will result in a list of contacts, and it must be determined whether some of these contacts correspond to potential new threats. In order to do this, the list of contacts is compared to a reference list compiled previously. The contacts that are in the new list but not in the old one are marked as potential targets that will require further investigation.

Comparison of two lists of contacts is complicated by several factors which influence the apparent location of the contacts: GPS positioning errors; slant-range to ground-range image corrections; and the accuracy of the AUV's inertial navigation system can result in overall positioning errors in the order of 10 meters between surveys. A solution to this inherent data inaccuracy is to use Rigid Data Association [10] between the different sets of contacts to be matched. This results in the rigid geometrical transformation that optimally maps one set to the other, accounting for positioning errors. After this geometrical correction is applied, contacts can be matched by greedy association [10]. Those contacts from the new set that remain unmatched after the association stage are, in effect, previously unseen contacts and will be flagged as potential threats.

The rigid Data Association procedure [10] requires a set of parameters describing the navigational system of the AUV to be set in advance. In MUSCLE's case navigational tolerances for position, orientation and scale were as follows: dt = 3m, dα = 0.115°, ds = 0.002. Non-recoverable error dmax was set to 10m.

5. RESULTS

During COLOSSUS 2 NURC collected over 20km^2 of sonar data, covering many different seabed types. This study concentrates on a small area of this total data set that includes a mixture of flat and rocky bottom types and areas of both high and low clutter. Multi-aspect views of the area were conducted both before (Day1) and after (Day2) deploying a number of control target shapes, enabling the performance of the change detection to be analyzed for different search directions or their combination.

Figure 2-left shows the coverage density (number of times a particular spot has been observed) and sonar tracks, before and after deployment of the targets.

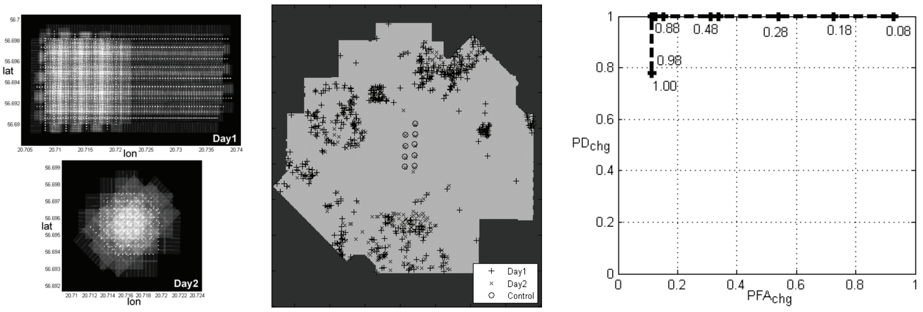

Fig.2: Left: Area coverage densities and sonar tracks for Day1 (top) and Day2 (bottom). Middle: contacts found on Days 1 and 2 in the intersecting area covered by both surveys. Right: ROC curve for the Change Detection procedure when comparing the unmatched contacts to the list of known control targets, using the "persistence" ratio as the varying parameter

The change detection can only be performed in areas that were covered on both days. For this, the intersecting region of the coverage maps for Days 1 and 2 was computed. Contacts located in that region were then associated and those that were deemed new events were compared against the list of control targets in order to determine the performance of the proposed change detection strategy. Figure 2-middle shows the intersecting region and the contacts found on the two days. On Day1 1845 contacts were found by the detector; on Day2, 371. These sets of contacts were associated by the proposed method, and those falling outside the intersecting region were excluded from further analysis.

Comparison of the newly found contacts within the intersecting region against the list of known deployed targets permits the performance of the proposed Change Detection procedure to be determined. The main factor regulating this performance was found to be the number of times a given contact is detected. Target-like objects will raise a detection almost every time they are observed (with a probability PD_{det}, close to 1) while non-target-like objects will only sporadically produce a false alarm, with doubtful objects falling somewhere in the middle. However, since some areas are covered by more sonar tracks than others, it is important to normalize the number of times a contact is detected by the number of times it has been observed, which is what we call the "persistence" or "significance" ratio. A significant new contact (meaning one that wasn't there before and is also target-like) should have a persistence ratio close to one, since they should appear as detections almost every time they are observed.

The persistence ratio can be used to form the ROC curve for the Change Detection procedure, which is shown in Figure 2-Right. From the curve we can conclude that detection of every new threat appears possible (PD_{chg}=1) while maintaining a false alarm rate of about PFA_{chg}=0.1. This translates on roughly 30 false alarms for the survey area (intersection region of Day1 and Day2), which covers 42000m^2. Although this number may be high for some applications, an operational implementation of Change Detection requires another additional stage for analyzing the potential new threats found. It is expected that the application of ATR algorithms such as those described in [8] will suffice to discard most of the false alarms found. Our future work will aim in that direction.

A question of particular interest is how different survey directions affect the results of the change detection procedure. Thus the lists of detected contacts were grouped by the orientation the legs were run, and then the change detection procedure was performed. On both days, tracks were run East to West and North to South. The results of the change detection procedure for the different combinations between search directions are presented in Figure 3.

Figure 3 shows that the change detection procedure is not very sensitive to using different survey directions on Days 1 and 2, although better results are indeed obtained when using the same directions for both days. The main factor responsible for the shape of the ROC curves seems to be the set of contacts found on Day2, which fix the PD_{chg} values in the curves.

In any case, what is clear from the comparison of any of the curves in Figure 3 with the one in Figure 2 is that using all the information available (all orientations on both days) offers the best possible change detection performance.

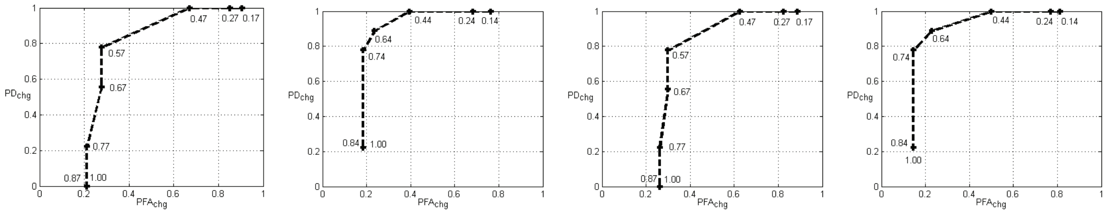

Fig.3: The Change Detection procedure applied to combinations of tracks from single search directions: (a) Day1 E-W versus Day2 E-W, (b) Day1 E-W versus Day2 N-S, (c) Day1 N-S versus Day2 E-W, (d) Day1 N-S versus Day2 N-S.

6. CONCLUSIONS

A simple and efficient approach to Change Detection has been presented in this paper. The approach consists of two main stages. First, the application of a basic target detector to the two surveys to be compared is used to obtain a list of contacts for each of them. Second, the lists of contacts are matched by rigid data association in order to find previously unseen ones.

The paper has demonstrated the importance of performing multi-coverage multi-orientation surveys in areas of higher seabed complexity. The persistence factor is an effective and convenient way of weighting the contributions of the different observations. Performing the surveys in a way that maximizes the potential persistence can greatly improve the change detection quality by reducing false alarms.

A third stage is envisaged where the list of new potential threats is filtered by an ATR algorithm [8][11]; this constitutes the topic of our current research work.

It is important to mention that the short time between the first and second surveys removed the potential for the environment affecting the results by covering or uncovering seabed objects. However the study proves that an AUV can successfully conduct this type of change detection operation without having to follow exactly the same survey patterns.

7. ACKNOWLEDGEMENTS

NURC would like to acknowledge the excellent support and assistance that was received from the Latvian Navy and MOD while gathering the data used in this paper.

REFERENCES

[1] **S. C. Truver**, “Mines and Underwater IEDs in US Ports and Waterways”, Naval War College Review, vol. 61, no. 1, 2008.

[2] **G. V. Jean**, “Improvised Explosive Devices: Could They Threaten U.S. Ports?”, National Defense, January 2008.

[3] **S. C. Truver**, “Mines, improvided explosives: a threat to global commerce?”, National Defense, April 2007.

[4] **European Community Shipowner’s Association,** “Press Conference 04-2, ECSA Annual Report 2003-2004”, Monday, 4 October 2004.

[5] **CRS Report for Congress**, “Maritime Security: Potential Terrorist Attacks and Protection Priorities”, January 9, 2007

[6] **Port of Rotterdam,** “Port Statistics”, 2007.

[7] **G. J. Dobeck, J. C. Hyland, L. Smedley**, “Automated detection/classification of seamines in sonar imagery”, Proceedings of SPIE-ISO, vol. 3079, pp. 90-110, 1997.

[8] **J. Groen, E. Coiras, J. Del Rio Vera, B. Evans**, “Echo classification features and processing for mine-hunting with autonomous underwater vehicle”, NURC-FR-2008-005, February 2008.

[9] **F. Baralli, B. Evans, E. Coiras, A. Bellettini**, “AUV Navigation Systems for MCM Operations”, NURC-FR-2007-007, May 2007.

[10] **E. Coiras, F. Baralli, B. Evans**, “Rigid Data Association for Shallow Water Surveys”, IET Radar, Sonar and Navigation, Vol.1, No.5, pp. 354-361, 2007.

[11] **E. Coiras, J. Groen, V. Myers, B. Evans, “**Estimation of 3D Shape from High-Resolution Sonar Imagery for Target Identification”, in Proceedings of the IoA DCUT 2007 conference, Edinburgh, 2007.