

Automated Change Detection Using Synthetic Aperture Sonar Imagery

Marlin Gendron, Maura Lohrenz and John Dubberley
Naval Research Laboratory Code 7440.1, Stennis Space Center, MS 39529

Abstract—When resurveying a seafloor area of interest during change detection operations, an automated method to match found bottom objects with objects detected in a previous survey allows the surveyor to quickly sort new objects from old. The change detection system developed at the Naval Research Laboratory contains modules for automatic object detection, feature matching using shadow outlining, scene matching using control-point matching, and visualization capabilities. This system was developed for sidescan sonar surveys using instrumentation such as the high-frequency Marine Sonic Technology sidescan sonar. In this paper, the authors describe modifications to the sidescan-based system required to perform change detection using Synthetic Aperture Sonar (SAS) bottom imagery.

Index Terms—Acoustic signal detection, Computer aided analysis, Real time systems, Sonar signal processing.

I. INTRODUCTION

SEAFLOOR change detection (figure 1) is accomplished by comparing high-resolution sidescan sonar imagery (SSI) collected sometime in the near past (“historical”) with recently collected SSI in an attempt to identify newly placed objects (“contacts”) that might pose a threat, such as explosive mines and Improvised Explosive Devices (IEDs). Any new contacts not successfully matched with historical ones are flagged for further investigation, typically using a Remotely Operated Vehicle (ROV) with a camera. This form of change detection, performed both manually and automatically, has been shown to greatly reduce the time and resources needed to determine if harbors and routes for ship traffic are safe [1].

The imagery typically used for change detection is high-frequency SSI with resolution of between 0.1-0.5 meters per pixel (m/pixel). For example, the Klein 5000 sonar operates at a frequency of 450 kHz and has an across-track resolution of 0.1 m/pixel. The high-frequency Marine Sonic Technology sidescan sonar, commonly used on the Remote Environmental Monitoring Unit (REMUS) Autonomous Underwater Vehicle (AUV), operates at 900 kHz with a resolution of 0.586 m/pixel. At these frequencies and resolutions, it is possible to detect and classify bottom objects as possible threats, but not to identify the specific type of threat.

Imagery can be generated from synthetic aperture sonar (SAS) with resolutions that can exceed 10 times that of SSI. The resulting detail might support the identification of specific objects [2] thus improving the ability to perform change

detection and threat neutralization. The higher resolution is achieved when SAS ensonifies the same position on the seabed with multiple pings, which are then coherently reorganized.

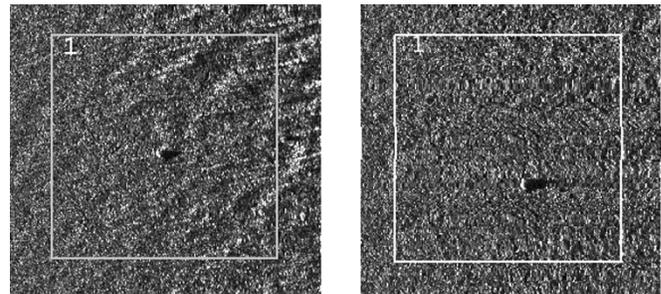


Fig. 1. Close up view of the same contact observed in the past (left) and present (right). Note the contact in the historical SSI (right) is not centered in the box. This is due to position error between historical SSI and recently collected SSI.

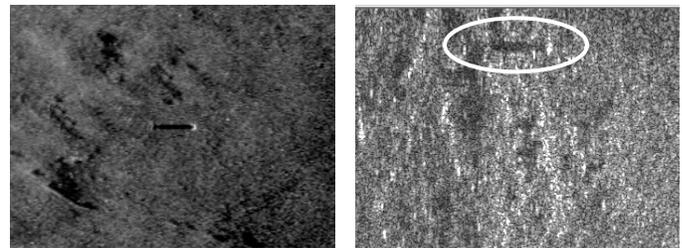


Fig. 2. SSI (left) and SAS (right) of the same area. The object circled in the SAS is the same as that centered in the SSI. Note the SAS imagery is much lighter than the SSI, and shadows are less evident. In addition, the SAS imagery appears to have considerably lower signal to noise ratio than the SSI, in this example.

In May 2009, researchers at the Naval Research Laboratory (NRL) participated in a homeland defense Limited Objective Experiment (LOE) with the Naval Oceanographic Office (NAVOCEANO) and the Naval Oceanography Mine Warfare Command (NOMWC) to look for mines in the Corpus Christi harbor [3]. A few months prior to the LOE, operators used REMUS AUVs equipped with sidescan sonar systems to pre-survey the harbor. At the start of the LOE, inert mineshares were placed in the ship channel to simulate a terrorist threat. During the LOE, operators resurveyed the harbor with both sidescan sonar (on REMUS) and SAS (on the Small Synthetic Aperture Minehunter AUV), shown in figure 2. NOMWC, NAVOCEANO and NRL then analyzed the SSI and SAS to find the inert mineshares.

This paper presents efforts to perform automated change detection between historical SSI (collected during the pre-survey) and the new SAS imagery (collected during the Corpus Christi LOE) using the authors' Automated Change Detection and Classification (ACDC) system. ACDC currently automates the change detection process with sophisticated algorithms operating on high-resolution SSI [4-9]. The authors first discuss briefly how their Automated Target Recognition (ATR) component of ACDC currently works with SSI. They describe how, for a proof-of-concept test, they processed SAS into a format that could be input directly into the NRL ATR, by downsampling SAS and stretching its intensity range to mimic SSI. They present preliminary results of the Corpus Christi LOE comparing how well the NRL ATR detected objects in the modified SAS imagery vs. SSI, and finally discuss how the ATR could be modified to operate on the original, full-resolution SAS imagery.

II. NRL AUTOMATED TARGET RECOGNITION (ATR)

Seafloor objects show up as bright spots ("brights") and shadows in SSI. If the shadow is on the side of the object furthest from nadir (the point directly below the sonar), then the object is proud on the seafloor and its height can be estimated from the length of its shadow and the height of the sonar above the seafloor. The authors' ATR algorithm detects mine-like objects, called "clutter", by locating brights with adjacent shadows in SSI. The algorithm rejects clutter that is outside a range of sizes consistent with known threats.

SAS data can be processed into imagery that resembles SSI, with brights and shadows representing bottom objects. Experimental SAS processing initially performed by the Naval Surface Warfare Center - Panama City (NSWC-PC) produced imagery that does not represent the full potential of SAS. As the NSWC-PC SAS algorithms improve, image quality is expected to improve. In the meantime, this imagery displayed shadows that were not as dark as those typically seen in SSI, due to the method by which the multiple SAS pings were correlated. The NRL ATR was designed to be triggered by sufficiently dark shadows with associated brights. As a preliminary proof-of-concept during the Corpus Christi LOE, NRL applied standard image processing techniques to the SAS imagery (to darken the shadows) prior to running their ATR. This was the first opportunity NRL had to work with the NSWC-PC SAS imagery. If this initial test were successful, then NRL could provide change detection results by the end of the LOE. If not, then in a future effort, NRL would investigate more sophisticated image processing techniques and modify the ATR to correctly handle SAS imagery.

The NRL ATR can operate in real-time as sonar data is collected and is the first component of the real-time version of ACDC (ACDC-RT). Due to real-time processing demands, the ATR algorithm relies on a patented geospatial bitmap (GB) technique. Simple bitmaps, with a depth of one bit per pixel, are binary structures in which bits are turned on (set = 1) or off (cleared = 0). The index of each bit is unique and denotes its position relative to other bits in the bitmap. The

authors extended this concept to construct GBs in which every bit represents an object at some unique geospatial location. A set bit indicates that some object of interest (in this case a bright or shadow pixel) exists at that location, accurate to within the resolution of the bitmap. A cleared bit indicates the absence of any object of interest at that location. Although a GB can be defined for an entire finite space, memory is only allocated (dynamically) when groups of spatially close bits are set, resulting in a compact data structure that supports very fast Boolean and morphological operations. Thus, the GB structure is well-suited for real-time imagery operations.

Across-track bright and shadow positions and lengths are stored in two GBs: one each for shadows and brights. Shadows and brights in a scan line are located by first adaptively obtaining a lower intensity threshold, i_{\min} , such that all samples of intensity less than i_{\min} are considered shadows. An upper intensity threshold, i_{\max} , is set such that all samples of intensity above i_{\max} are considered brights.

An appropriate gamma shift converts image intensities to fit a normal distribution, such that i_{\min} and i_{\max} are set to the quartiles of the shifted (normal) distribution. After i_{\min} and i_{\max} have been determined for scan lines with maximum intensity value > 128 , the port and starboard halves of the scan lines are processed separately. Each half of the scan line can be represented by a vector, X , of length N . The following method is used to process shadows and brights for the starboard side; the port side is processed similarly.

Two GBs of size $1 \times N$ are created, one for shadows and one for brights. A different gamma adjustment, γ , based on an error approximation of the side-scan sonar parameters, is computed for position x within X , as shown in (1).

$$\gamma = e^{-\beta x/N} \quad (1)$$

β is based on the sonar parameters, such as time-varying gain. As β approaches infinity, the gamma correction approaches 0 over a greater range of X and therefore has less affect on intensity thresholds.

The bright and shadow thresholds $I_{\min}(x)$ and $I_{\max}(x)$ are defined in (2) and (3). All pixels with intensity values above $I_{\max}(x)$ are considered brights, all with intensities below $I_{\min}(x)$ are considered shadows, and the corresponding bits in the bright and shadow GBs are set.

$$I_{\min}(x) = i_{\min}(1-\gamma) \quad (2)$$

$$I_{\max}(x) = i_{\max}(1+\gamma) \quad (3)$$

Figure 3 illustrates how the intensity thresholds vary over x for a given γ . For example, the closer a pixel is to the center of the scan, known as nadir ($x = 0$), the greater its intensity must be to be detected as a bright [7], and the lower its intensity must be to be detected as a shadow. It is interesting to note that the two threshold curves do not diverge from their respective asymptotes (i_{\min} and i_{\max}) at the same rate as they approach nadir. This is by design, because shadows are more detectable than brights in SSI [8] – [11]. In other words, a single shadow threshold value (i_{\min}) suffices for more values of x than a single bright threshold value (i_{\max}).

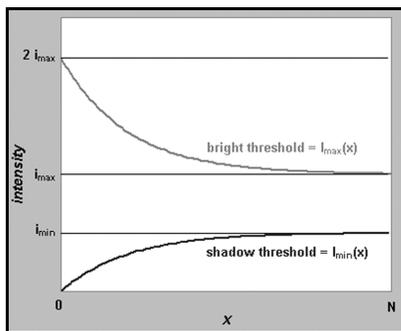


Fig. 3. Intensity thresholds for brights and shadows.

Finally, the bright and shadow geospatial bitmaps are examined from the edges of the scan lines toward the center (nadir) to detect runs of shadows followed by runs of brights. A circular lookup table is created to “window” several scan lines at a time. This lookup table is populated with positions and run-lengths of shadows and brights. The window information is used to determine if a series of scan-line detections comprise an object. Shadow length is one component in determining the object’s height, which is used to help determine whether the object is mine-like.

III. ATR TEST METHODS

A. Pre-processing SAS

To prepare the SAS imagery for input to the ATR, as a proof-of-concept test, NRL first converted the original 16-bit SAS imagery to 8-bit. The original imagery is stored in NAVOCEANO Unified Sonar Image Processing System (UNISIPS) files, in which one sample value occupies two bytes (with an intensity range of 0-65535). Each 2-byte sample is converted to 1 byte (with an intensity range of 0-255) by computing the actual intensity range of the original image and rescaling it to 0-255. NRL then downsampled the imagery from the original 0.0235 meters per pixel (m/pixel) to 0.4700 m/pixel by removing an appropriate number of pixels. This approximates the REMUS imagery resolution (0.586 m/pixel). Finally, to darken the shadows in the SAS imagery, NRL reduced the brightness by 10% and increased contrast by 33%, which produced a range of intensities that approached that of the REMUS data (figure 4).

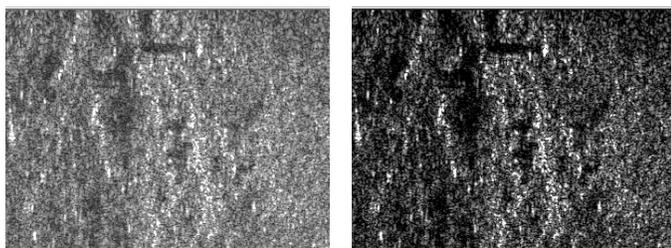


Fig. 4. Downsampled SAS before (left) and after (right) darkening to emphasize shadows (decreased brightness by 10%; increased contrast by 33%).

B. Performing ATR on SSI and SAS

NRL selected a section of the Corpus Christi harbor where both SAS imagery and SSI were collected during the LOE and

the locations of 690 bottom objects had been determined manually by Navy operators. NRL ran their ATR algorithm on both sets of imagery for this region, and determined whether or not each detection was correct (i.e., within 2 m of a manually called object). The total number of correct detections and false detections were tabulated for each sonar type. The number of missed detections was assumed to be 690 minus the number of correct detections.

IV. RESULTS

Figure 5 and table 1 present ATR results using SSI and SAS imagery from the 2009 Corpus Christi LOE. Detection was significantly better using SSI than SAS in this experiment. Of 690 seafloor objects identified manually by Navy operators, the NRL ATR correctly detected 601 (87%) using SSI and 182 (26%) using SAS. The ATR false detection rate (type I error) was 0.06 for SSI and 0.20 for SAS, and the missed detection rate (type II error) was 0.12 for SSI and 0.59 for SAS.

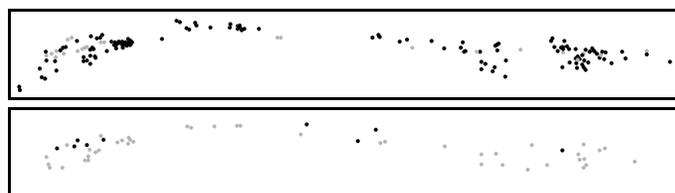


Figure 5. Comparison of correct detections (top) and false detections (bottom) for ATR performance with REMUS (black) and SAS (gray) imagery, over the same test area.

TABLE I
ATR PERFORMANCE WITH SSI vs. SAS DURING CORPUS CHRISTI LOE.

	SSI		SAS	
	Present	Not Present	Present	Not Present
Detected	601	46	182	173
Not detected	89	0	508	0
Total	690	46	690	173

	SSI	SAS
α (Type I error)	0.06	0.20
β (Type II error)	0.12	0.59
$1-\beta$ (Power of test)	0.88	0.41

$$\alpha = \text{False positive rate} = \text{\#false detections} / \text{total \#cases}$$

$$\beta = \text{False negative rate} = \text{\#missed detections} / \text{total \#cases}$$

V. CONCLUSIONS

The authors’ objective was to use ACDC to perform automated change detection between SAS imagery (collected during the Corpus Christi LOE) and historical SSI (collected months prior to the LOE). The authors first needed to determine whether the NRL ATR (the first component of ACDC) would perform satisfactorily on SAS data. To test this in time for the Corpus Christi LOE, NRL reformatted the SAS data to approximate the resolution and intensity range of SSI for easy input to the ATR.

Unfortunately, the ATR performed much worse with the pre-processed SAS imagery than with SSI, by not detecting most (74%) of the bottom objects called manually by operators and reporting a very high number of false detections. While it is possible that this poor performance was due to information lost during preprocessing (i.e., downsampling the SAS data and scaling its image intensities), the authors believe ATR results would be greatly improved with better quality SAS imagery. The SAS imagery produced by NSW-PC during the LOE is experimental and has not yet reached its full potential (i.e., sharper imagery with greater detail, less noise, darker shadows and brighter brights).

Given the poor ATR performance with this SAS imagery, NRL decided it was not feasible to perform automated change detection between historical SSI and new SAS imagery for this experiment. However, NOMWC successfully performed manual change detection by visually comparing historical SSI with the newly collected SAS imagery.

As NSW-PC continues to improve their SAS processing techniques, NRL hopes to have access to better quality SAS imagery. At that point, NRL plans to reengineer the ATR to better detect objects in SAS by modifying our gamma shift to handle softer shadows, and increasing the amount of data being buffered so the ATR accommodates the full resolution of SAS. Once the ATR has successfully detected objects in SAS, NAVOCEANO and NRL can start testing ACDC with SAS imagery.

ACKNOWLEDGEMENT

This work was sponsored under Program Element 602435N by the NRL 6.2 Base Program. The authors thank Mr. James

Hammack at NAVOCEANO for his technical support during this project. We also thank Dr. Gerald Dobeck and Mr. Jose Fernandez at NSW-PC for sharing the SAS data with us.

REFERENCES

- [1] Lingsch, W.C. and S.C. Lingsch (2001). Sonar Image Change Detection as a Mine Counter Measures Tactical Tool (abstract). *Proceedings of the Shallow Survey 2001 Conference*, Sidney, Australia.
- [2] Willcox, S., J. Bondaryk, K. Streitlien, C. Emblen, and J. Morrison (2007). *A Bluefin-12 based system solution for the US Navy's Surface Mine Counter-Measures Unmanned Underwater Vehicle Program: Increment 2 (SMCM/UUV-2)*. Bluefin Robotics Corp. Technical Paper. http://bluefin2.vh.shore.net/press_releases.htm. May 3.
- [3] Lammons, G. (2009). *Oceanographers Operate Unmanned Vehicles to Find Mines at LOE*. www.navy.mil Story #NNS090522-32. May 22.
- [4] Gendron, M., G. Layne, C. Gautre, J. Hammack and C. Martin (2006). The Automated Change Detection and Classification – Real-time (ACDC-RT) System. *Proceedings of the Undersea Defense Technology Pacific Conference*. Dec. 6.
- [5] Gendron, M. and M. Lohrenz (2007). The Automated Change Detection and Classification Real-time (ACDC-RT) System. *Proceedings of the IEEE Oceans 2007- Europe Conference*. Aberdeen, Scotland. June 18.
- [6] Gendron, M., M. Lohrenz, G. Layne and J. Ioup (2004). Automatic Change Detection and Classification (ACDC) System. *Proceedings of the Sixth International Symposium on Technology and the Mine Problem*. Naval Postgraduate School, Monterey, CA. May 9-13.
- [7] Gendron, M.L., M.C. Lohrenz, L. M. Riedlinger, (2009). *Computer-implemented real-time side scan sonar imagery; i.e., unified sonar image processing system format data, processing method, involves displaying historical side scan sonar imagery converted to area-based tiles*. U.S. Patent 2009016161-A1.
- [8] Lohrenz, M., M. Gendron and G. Layne (2004). Automatic Change Detection and Classification (ACDC) System in support of homeland security. *Proceedings of the 2004 IEEE OES Homeland Security Technology Workshop*. Valley Forge, PA. Dec. 6-8.
- [9] Lohrenz, M., M. Gendron and G. Layne (2005). *Automated Change Detection and Classification (ACDC) System*. 2005 NRL Review, pp.155-157.