Automated Change Detection for Synthetic Aperture Sonar

Tesfaye G-Michael*[a], Bradley Marchand[a], J Derek Tucker[a], Daniel D Sternlicht[a], Timothy M Marston[b], Mahmood R. Azimi-Sadjadi[c]

[a]The Naval Surface Warfare Center Panama City Division Panama City, FL 32407-7001 USA; [b]The University of Washington Applied Physics Laboratory [c] The Colorado State University

ABSTRACT

In this paper, an automated change detection technique is presented that compares new and historical seafloor images created with sidescan synthetic aperture sonar (SAS) for changes occurring over time. The method consists of a four stage process: a coarse navigational alignment; fine-scale co-registration using the scale invariant feature transform (SIFT) algorithm to match features between overlapping images; sub-pixel co-registration to improves phase coherence; and finally, change detection utilizing canonical correlation analysis (CCA). The method was tested using data collected with a high-frequency SAS in a sandy shallow-water environment. By using precise co-registration tools and change detection algorithms, it is shown that the coherent nature of the SAS data can be exploited and utilized in this environment over time scales ranging from hours through several days.

Keywords: automated change detection, canonical correlation analysis, coherent change detection, co-registration, synthetic aperture sonar, SAS

1. INTRODUCTION

Detecting areas of change in multiple images of the same scene captured at different times is of interest to vastly different disciplines. Applications of change detection include: remotely sensed satellite imagery;[1,2] synthetic aperture radar (SAR);[3] medical imaging;[4–6] industry applications;[7,8] and driver assistance systems.[9,10] For sonar and radar systems, change detection is the process by which regions of interest are identified through the comparison of current data with historical.

In image-based change detection, entire image portions are compared to identify regions of change. Myers, et al.,[11] used several change detection methods to calculate overall degree of similarity between two sidescan sonar images. Typically, change detection is applied to situations where the same area is to be repeatedly monitored, such as surveys for port and harbor security. This change detection technique would increase detection capability and reduce the workload of a human operator tasked with identifying anomalies in sidescan sonar images.

Techniques for image-correlation automated change detection (ACD) have been under development by the SAR community since at least the 1990s,[3] and procedures to fuse scene changes derived from segmented features with pixel or parcel based change maps have recently been explored.[12] In the sonar research community, contact correlation ACD approaches were developed first.[13] This was due to a number of factors, such as challenges in the trajectory control, accurate georegistration, the complexity of the propagation medium, and the radiometric inconsistencies of conventional sidescan sonars. Only in recent years, with the advent of synthetic aperture sonars (SAS) and actively navigated platforms, have image correlation methods been shown to be feasible.[14–17]

The inspiration for the current work comes from an extension of the coherent-based change detection results using canonical correlation analysis (CCA) described by Azimi-Sadjadi and Srinivasan,[18] G-Michael and Tucker[15] and Sternlicht and G-Michael,[19] where the preliminary studies were performed on simulated SAR and SAS imagery. The motivation behind CCA comes from recent success using it as a target detection scheme in SAS imagery.[20] In this paper, we develop a complete ACD algorithm that combines the use of the Scale Invariant Feature Transform (SIFT) algorithm[21,22] as a co-registration tool with the change detection method of using CCA. In CCA, the correlation is formed between two sonar images from repeat passes over the ocean floor. From

*tesfaye.g-michael@navy.mil — Phone: (850) 235 - 5295 — Fax: (850) 235 - 5374

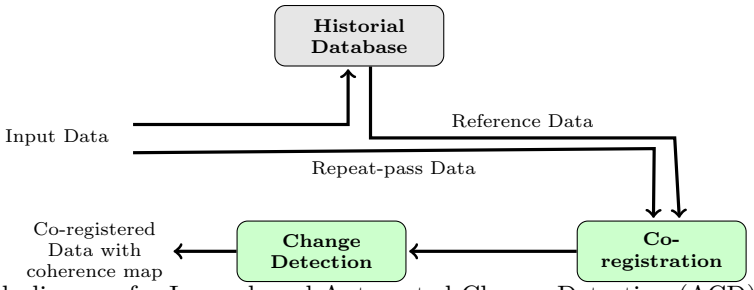

Figure 1: Block diagram for Image-based Automated Change Detection (ACD) processing.

this correlation, coherence (or incoherence) can be measured and used to determine what has changed between the two sonar images.

Our methods are tested on SAS imagery collected with the high-frequency projector of the Small Synthetic Aperture Mine-hunter (SSAM) system developed by the Naval Surface Warfare Center Panama City Division and the Applied Research Laboratory, Pennsylvania State University.[23] The SSAM system includes an advanced sidescan sonar that is mounted on a Woods Hole Oceanographic Institution REMUS 600 autonomous undersea vehicle (AUV). It is able to achieve fine-detail high-resolution SAS imaging of the seabed using a high-frequency (>100 kHz) projector.

This paper is organized as follows: Section 2 presents an overview of the proposed image-based ACD process. Section 3 provides detailed explanation of our co-registration method. Section 4 describes the CCA algorithm and how they measure change. Section 5 presents performance results of our co-registration and change detection methods. Finally, Section 6 provides conclusions, discussion, and directions for future research.

2. AN OVERVIEW OF ACD PROCESS

Image-based change detection is carried out through the comparison of reference and repeat-pass images that correspond to the same geographical location, but are collected at different times. The automation of this task involves three major components: historical data retrieval, reference and repeat-pass co-registration, and detection of changes. A block diagram of ACD processing is illustrated in Fig 1.

A historical database comprised of SAS images with their associated geographical locations provides the baseline or reference. Given a specific repeat-pass image, images from this database that roughly correspond to the same geographical location are retrieved for use in the following co-registration and change detection steps. All image geographical information for both the database and repeat-pass images are collected from the REMUS 600 on-board navigation system and is accurate up to the fidelity of this system, which is comprised of a Doppler velocity log and GPS-aided inertial navigation system.

3. CO-REGISTRATION

Image co-registration is a key step in the ACD process. It provides the image registration necessary for image comparison and change detection. For our co-registration method, we propose a three stage process: navigation alignment, fine-scale co-registration, and sub-pixel co-registration.

3.1 Navigation Alignment

After the identification and loading of the historical data corresponding to the repeat-pass image, a purely navigational coarse alignment is our first step toward co-registration of the images. The process involves using the nominal paths for the reference and repeat-pass images. The nominal path for an image is the straight line path taken from the image starting location to the end. For the repeat-pass image we construct a linear mapping from latitude and longitude sample location to their relative pixel location by solving the following system of equations:

$$\begin{bmatrix} 1 \\ m \\ m \end{bmatrix} = \begin{bmatrix} \boldsymbol{x}_1^\intercal & 1 \\ \boldsymbol{x}_2^\intercal & 1 \\ \boldsymbol{x}_3^\intercal & 1 \end{bmatrix} \boldsymbol{c}_1 \tag{1}$$

$$\begin{bmatrix} 1 \\ 1 \\ n \end{bmatrix} = \begin{bmatrix} \boldsymbol{x}_1^\intercal & 1 \\ \boldsymbol{x}_2^\intercal & 1 \\ \boldsymbol{x}_3^\intercal & 1 \end{bmatrix} \boldsymbol{c}_2 \tag{2}$$

where the size of the repeat-pass data is $m \times n$, $\{\boldsymbol{x}_1, \boldsymbol{x}_2, \boldsymbol{x}_3\}$ are the latitude and longitude of the endpoints of the repeat-pass data, and $\{\boldsymbol{c}_1, \boldsymbol{c}_2\}$ are the linear mapping coefficients that we are solving for. Using this mapping we can calculate the reference data pixel values relative to the repeat-pass data, which allows interpolation of the repeat-pass data to the reference data relative to pixel location. This projects the repeat-pass data on to the reference, thereby removing most navigation differences between reference and repeat-pass images.

3.2 Fine-Scale Co-Registration

For fine-scale co-registration we apply the scale invariant feature transform (SIFT) algorithm[21,22] to the amplitude image data. The SIFT algorithm detects and extracts distinct features from images which can be used to perform matching between overlapping images. The SIFT features computed for each keypoint are gradient magnitudes and orientations that are sampled around the keypoint location. In order to achieve rotation invariance, gradient orientations are rotated relative to the keypoint orientation. SIFT keypoints and features are computed for both navigation-aligned reference and repeat-pass data. Keypoint matching between reference and repeat-pass images is done by computing minimum Euclidean distance between SIFT features. In order to reject weak keypoints, matches where the second closest neighbor is within 1.5 times the closest match are rejected.

Differencing the matched keypoints typically results in a collection of noisy estimates for translational shifts. In order to approximate the true shift, we first remove outliers by iteratively removing shift estimates that are outside three standard deviations from the mean.

Once outliers have been removed, our final shift estimate is computed as a weighted average of the remaining shift estimates. Specifically, letting $\{\boldsymbol{s}_1, \boldsymbol{s}_2, \ldots, \boldsymbol{s}_N\}$ be our remaining collection of shift estimates and c_{ij} be the Euclidean distance from $\boldsymbol{s}_i$ to $\boldsymbol{s}_j$. We compute the weights, $\{w_1, w_2, \ldots, w_N\}$, as

$$w_i = \frac{1}{\min_{j \neq i} c_{ij}}.$$

For weights where $\min_{j \neq i} c_{ij}$ is zero, we set the weight value to twice the maximum of the defined weights. The weights are then normalized,

$$w_i = \frac{w_i}{\sum_{j=1}^{N} w_j},$$

and the image shift is estimated as

$$\boldsymbol{s} = \sum_{i=1}^{N} w_i \cdot \boldsymbol{s}_i. \tag{3}$$

An alternative method for estimating the image shift is to fit a Gaussian distribution to the remaining shift estimates. The shift is then approximated as the mean of the fitted Gaussian distribution. This method works well when there is a large number of shift estimates to use for fitting. However, we have found that with more difficult environments and longer elapsed times between passes, the quality of the resulting estimate is not as good as the estimate obtained with the weighted average.

With the fine co-registration shift offsets determined, both reference and repeat-pass images are cropped, such that only common pixel locations for the two images are passed to the change detection algorithm.

3.3 Sub-Pixel Co-Registration

Following fine-scale co-registration the average scene displacement is ideally on the order of a single pixel. However, local registration error can be significantly higher than the average, particularly if the deviation between platform trajectories is significant or if rotation is present. This spatially-variant mismatching can have a very large (and negative) effect on coherence estimates between scenes. The problem is similar to that encountered when co-registering wide-baseline interferometric channels.[24,25] Our sub-pixel co-registration works as an adaptive correlation technique, where the offset between images is a function of the $X-Y$ location within the image. It uses a quality metric to estimate whether coherent or intensity based correlation will result in optimal scene shift measurements. It works in such a way that, if the average correlation values computed over the overlapping portions of the projection of the sonar beam on the sea floor dip below a threshold, then intensity-based correlation values are calculated and a smoothness metric is evaluated for the best shift estimate results. The method which results in the best shift estimate is selected. A neighborhood of pixels with a fixed block size is chosen to allow for large offsets between images, while still maintaining a reasonably high correlation and the ability to generate a spatially variant solution. The resolution of the computed interferogram (a phase difference map implemented using the product of the reference image matrix and the complex conjugate of the co-registered repeat-pass matrix) and the resulting inter-scene coherence map can be improved by sub-pixel co-registration, because all the points in the corresponding neighborhoods overlap. Once the images have been aligned, there is a simple, one-to-one linear mapping of pixels between scenes, which is used to flatten the interferogram and improve the correlation between scenes.

Specifically, sub-pixel co-registration involves two steps: warping the complex values in the scene to align pixels on a local level, and flattening the interferogram such that the phase component due to the flat-earth phase trend is removed. In the current examples, sub-pixel co-registration is accomplished via the steps shown in Algorithm 1.

Algorithm 1 Sub-pixel co-registration

Step 1. Candidate scenes are autofocused
Step 2. Large-neighborhood (e.g. 25×25 pixel), non-overlapping patch correlation and parabolic filtering is used to estimate local patch alignment errors to pixel level
Step 3. Using an iterative, weighted least-squares approach, the X and Y shifts are independently fit to an appropriate dictionary of functions
Step 4. A two-dimensional warping function is defined using the coefficients calculated in step 3 and applied to the matching image
Step 5. Small-neighborhood (10×10 pixel) patch-correlation is used to generate an interferogram
Step 6. A weighted minimum-discontinuity unwrapping algorithm is applied to the interferogram
Step 7. A dictionary of functions defined by the estimated phase resulting from navigational deviations between repeat passes is fit using an iterative least-squares approach to the unwrapped phase
Step 8. An interferometric flattening function is generated using the calculated dictionary coefficients
Step 9. The flattening function is used to calculate a phase which is subtracted from the matching image

The full details of some of these steps are fairly involved and lie beyond the scope of this paper; however, it is similar to the algorithm described in Saebo, et al.[24] Particular steps merit additional attention. In some situations one can observe navigation-induced focusing defects, which can significantly degrade correlation between neighborhoods, even if alignment was fairly good. A non-parametric method similar to that used in multi-squint processing,[26] but highly local and modified to work with wide-beam/wide-band systems, was used to sharpen the focus of the beamformed scenes in step 1. The correlation coefficients calculated in steps 2 and 5 can be used for inter-scene coherence estimates. A qualitative comparison of the correlation coefficients after warping and phase flattening indicates that the primary enhancement in coherence occurs after the warping step. Typically, unless the baseline deviation is significant enough to produce strong local curvature in the phase flattening function, the phase correction is approximately planar and has minimal effect on the calculated coherence. Step 7 results in a heave and sway solution relating the difference in vehicle trajectory between passes, and in conjunction with the interferogram can be used to make repeat-pass bathymetry estimates.

4. CHANGE DETECTION

For change detection we utilize Canonical Correlation Analysis (CCA).[15] CCA is a multivariate statistical method[27] that determines the linear dependence (or coherence) between two data channels, which in this case correspond to reference and repeat-pass SAS images. The CCA method not only determines the amount of dependence (or independence) between two data channels but also extracts a subset of most coherent change features from the two channels. This is accomplished by linearly mapping the two data channels to their canonical coordinates, the first few of which capture maximal inter-channel coherence.

Let us assume that the two-channel complex data vectors are represented by $\boldsymbol{x} \in \mathbb{C}^n$ and $\boldsymbol{y} \in \mathbb{C}^m$, then the canonical correlations and coordinates are obtained by finding the singular values and vectors of the *coherence matrix*,[27] which is defined as

$$C = E[(R_{xx}^{-1/2}\boldsymbol{x})(R_{yy}^{-1/2}\boldsymbol{y})^H],$$

where $(\cdot)^H$ is the Hermitian operation, $E[\cdot]$ is the expectation operation, and R_{xx} and R_{yy} are the covariance matrices of $\boldsymbol{x}$ and $\boldsymbol{y}$ channels, respectively. The canonical coordinates of $\boldsymbol{x}$ and $\boldsymbol{y}$ are defined as

$$\begin{aligned} \boldsymbol{u} &= \mathrm{F}^H R_{xx}^{-1/2}\boldsymbol{x} \\ \boldsymbol{v} &= \mathrm{G}^H R_{yy}^{-1/2}\boldsymbol{y} \end{aligned}$$

where F and G are the mapping matrices containing the left and right singular vectors of the coherence matrix C, respectively. The canonical coordinate vectors $\boldsymbol{u}$ and $\boldsymbol{v}$ share the diagonal cross-covariance matrix

$$\mathrm{K} = R_{uv} = \mathrm{F}^H C \mathrm{G} \tag{4}$$

known as the *canonical correlation* matrix of canonical correlations $k_i, i \in [1, m]$, which are the singular values of matrix C.

Now, the linear dependence between $\boldsymbol{x}$ and $\boldsymbol{y}$ channels can be measured by

$$L = \prod_{i=1}^{N}(1 - k_i^2) \qquad 0 \leq L \leq 1.$$

This measure takes the value 0 if there is perfect linear dependence between $\boldsymbol{x}$ and $\boldsymbol{y}$; while $L = 1$ corresponds to the case when $\boldsymbol{x}$ and $\boldsymbol{y}$ are linearly independent. The *coherence* measure between $\boldsymbol{x}$ and $\boldsymbol{y}$ is given by

$$H = 1 - L \quad 0 \leq H \leq 1. \tag{5}$$

That is, the channels $\boldsymbol{x}$ and $\boldsymbol{y}$ are perfectly coherent if $H = 1$, and non-coherent if $H = 0$.

Though the dominant canonical coordinates capture maximal coherence, the subdominant ones maximize the divergence between the two channels $\boldsymbol{x}$ and $\boldsymbol{y}$.[15,18,28] In other words, the subdominant canonical coordinates capture most of the coherent change features between the two channels with a minimum dimensional feature set. The coherent change information between canonical coordinates $\boldsymbol{v}$ and $\boldsymbol{u}$ can be calculated using the residual, $\boldsymbol{v} - \mathrm{K}\boldsymbol{u}$, owing to the fact that $\mathrm{K}\boldsymbol{u}$ brings the most coherent information that the second channel $\boldsymbol{y}$ carries about the first channel $\boldsymbol{x}$. Now that the subdominant canonical coordinates capture the coherent change information (under the constraint of maximum rate) retaining the last few components of the difference vector $\boldsymbol{v} - \mathrm{K}\boldsymbol{u}$, is sufficient. This can be performed mathematically by using $I_p(\boldsymbol{v} - \mathrm{K}\boldsymbol{u})$, where I_p is the $m \times m$ diagonal matrix with the first $m - p$ diagonal entries being zero and the rest being unity. Moreover, we can convert the difference vector back to the image space from the canonical space using,

$$\boldsymbol{d} = R_{yy}^{1/2}\mathrm{G}I_p(\boldsymbol{v} - \mathrm{K}\boldsymbol{u}). \tag{6}$$

For a more complete description of the CCA change detection method for SAS imagery the reader is referred to.[15]

		Shift Error (Pixel)					
	Noise	1 <	20 <	50 <	100 <	150 <	200 <
set 1	σ	0%	0%	0%	0%	0%	0%
	$1.5 \cdot \sigma$	1%	0%	0%	0%	0%	0%
	$2.0 \cdot \sigma$	16%	8%	6%	6%	6%	6%
set 2	σ	0%	0%	0%	0%	0%	0%
	$1.5 \cdot \sigma$	0%	0%	0%	0%	0%	0%
	$2.0 \cdot \sigma$	31%	17%	16%	15%	14%	13%

Table 1: Percentage of images with estimated shift greater than true shift values for a feature rich environment (set 1) and a feature poor environment (set 2).

5. DATA AND RESULTS

To demonstrate the feasibility of our ACD approach we operated on SAS data collected with the high-frequency aperture of the SSAM system for a variety of measurement scenarios. We start with a qualitative demonstration of our approaches to co-registration and change detection, followed by a quantification of this performance by experiments that measure performance of both co-registration and change detection.

In Fig. 2 we show a qualitative, example of the result from the automated co-registration technique described in Section 3 applied in a flat sandy region. Typically co-registration in such regions is more difficult due to temporal changes of the base image structure. In this case, with a short time latency of 2 hours between passes, successful co-registration was possible. We were able to achieve high precision automated co-registration in this sandy environment for time latencies of up to 72 hours, then with more difficulty for longer time periods.

The result of applying the CCA change detection algorithm to the co-registered images shown in Fig. 2 can be seen in Fig. 3. In this figure the new object is clearly identified by the low-level region of the coherence map. Overall, we see a high level of coherence through out the center of this example and a low level at near and far range.

5.1 Robustness of Co-Registration Algorithm

In order to quantify the robustness of our co-registration algorithm, we constructed an evaluation set from the reference pass for data sets 1 and 2 by adding noise and shifting by a known quantity. The exact error of the co-registration algorithm can then be calculated as the Euclidean distance between the known image shift and the shift estimated by co-registration. The noise added was zero mean complex Gaussian. Three levels of noise were considered with variance levels of 1.0, 1.5, and 2.0 times the variance of the data. Each image was shifted down and right by 20, 50, 100, 150, and 200 pixels. Table 1 lists the percentage of images with estimated shift greater than the true shift values. Almost half of all the errors are less than 20 pixels in size, while most of the other errors are greater then 200 pixels in size. In general, the size of the shift difference does not impact co-registration performance for the shifts considered. For low and moderate noise, co-registration works well with typically less than a pixel of distance error in co-registration. For high noise, co-registration performs well for some images. When it does fail the estimated shift is vastly different from the true shift. These extremely large co-registration failures can likely be detected and flagged for operator intervention. The percentages in Table 1 also indicates the relative difficulty posed by the two environments. For the all-sand environment (data set 2), twice as many large co-registration failures occur at high noise, compared to the feature rich environment (data set 1).

5.2 Robustness of Change Detection Algorithm

In order to quantify the robustness of our change detection algorithm, we performed a statistical analysis on the measured temporal coherence of seafloor backscatter. A detailed comparison of correlation statistics is carried out based on computation of the CCA coherence measure, H(Eq. 5) over data sets collected at site 2. The duration of the data collection was approximately 72 hours over this predominantly flat sand environment. In Fig. 4, a box plot of the coherence distribution for the collected images in set 1 is shown. The coherence was

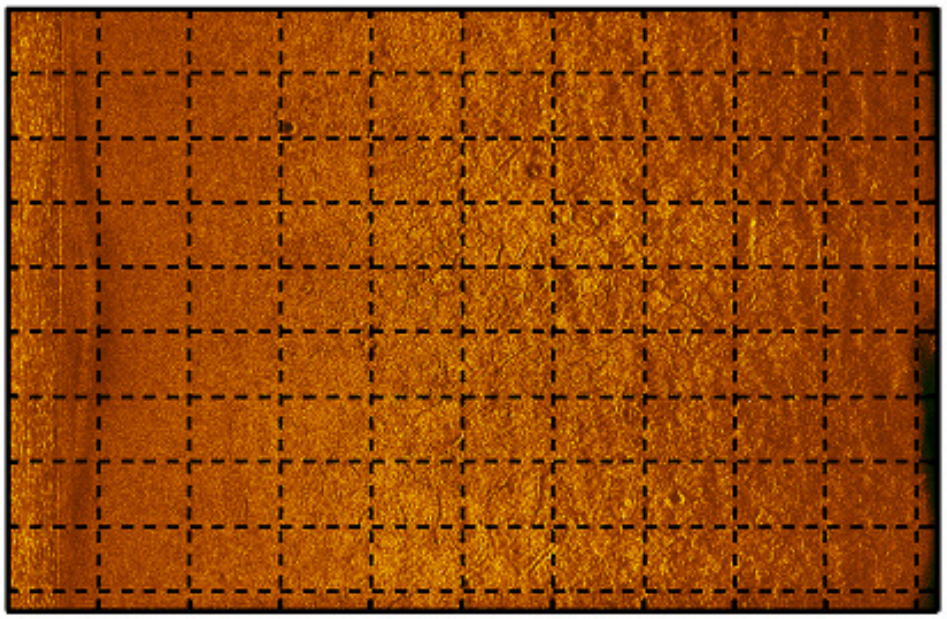

(a) Reference

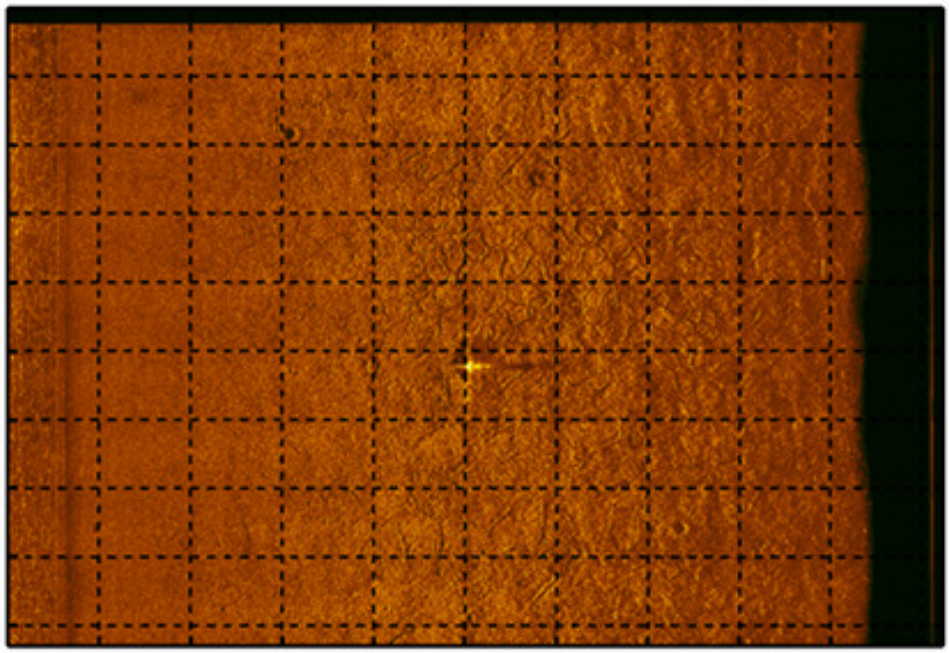

(b) Repeat pass

Figure 2: Automated co-registration results collected from data set 2. A new object is present in the center of the repeat-pass image (b).

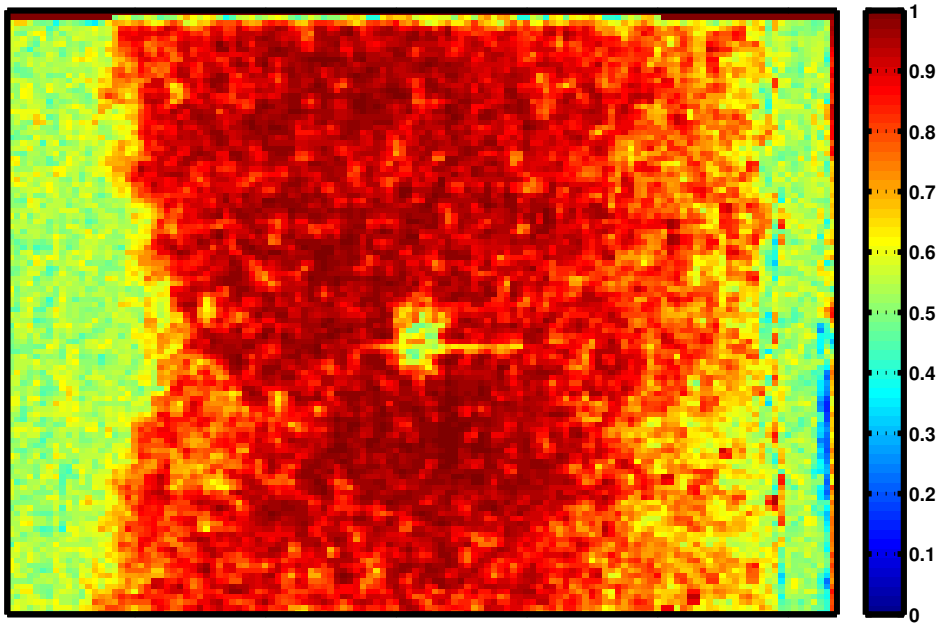

Figure 3: Coherence maps for co-registered images shown in Fig. 2 using CCA (a). Both maps range in value from 0 to 1, with 1 or red indicating high coherence and 0 or blue indicating low coherence.

computed between reference and repeat-pass data sets that are co-registered using the automated technique. The abscissae in Fig. 4 indicate the time interval of data collection for each repeat-pass leg. The boxes in these figures represent the middle 50% of the coherence distribution, with upper and lower hinges showing the 75^{th} and 25^{th} percentiles. The whiskers are extended to 1.5 times the difference between these percentiles (the inter-quartile range), showing that the majority of data falls between the ends of the whiskers. The horizontal red line through each box represents the median of the data within the box. The red crosses outside of the whiskers represent individual outliers.

As the correlations vary considerably from one run to another over time, the box plot allows us to better visualize the overall trend of coherence as time elapses. We observe a decaying of coherence over time which corresponds to temporal decorrelation. The first five repeat-passes were collected along $1-$hour intervals. The sixth and the seventh repeat-passes were collected approximately 72 hours later.

We determined the percentage of strong correlations and defined strong coherence as a value greater than 65%. The percentage of strong correlations should be regarded as an indication of the ability to perform image-based change detection. From Fig. 4 we see that the data is highly coherent within the first few hours with the exception of the data from the first hour where the vehicle altitude changed significantly between the reference and repeat-pass data. Removing the portion of the image associated with the water column prior to running the data through the change detection algorithm should mitigate this problem. Overall, the general decay of the coherence over time generally agrees with the observations of Lyons and Brown,[29] who quantify the temporal decorrelation of seafloor roughness and estimate scene decorrelation by calculating the complex correlation coefficient for pairs of SAS images generated using data collected with a rail-based system. The last two datasets in figures 4 represent data collection approximately 72 hours after the reference pass. They exhibit a much lower coherence, indicating that this sandy environment, lacking stable features, presents a challenge for change detection over large time scales.

6. CONCLUSIONS AND DISCUSSION

In this paper, we presented and demonstrated a new, robust, and precise seabed change detection method for synthetic aperture sonar that automatically compares seafloor images and detects changes over time. Our

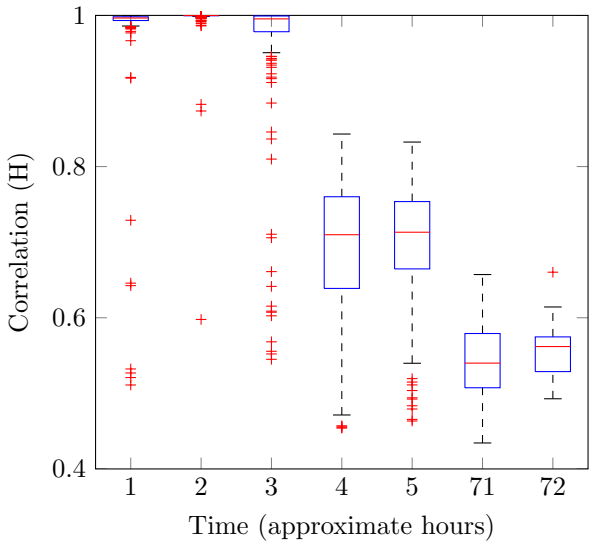

Figure 4: Coherence box plot. Sample box plots of the measured coherence of seafloor acoustic backscatter, H, over time. Seven repeat-pass data sets are compared to the reference data set with latencies ranging from 1 hour to 72 hours.

study primarily focused on the automation of the scene co-registration and change detection methods. The ACD method consists of a four stage process: 1) Coarse navigational alignment that relates and approximates the pixel location of reference and repeat-pass datasets; 2) Fine-scale co-registration using the SIFT algorithm to match features between overlapping datasets; 3) Sub-pixel co-registration that improves the coherent change detection process; 4) Coherent change detection utilizing the CCA and TIS algorithms to detect changes. The ACD method was tested using data collected with a high-frequency synthetic aperture sonar in a sandy shallow-water environment, subject to currents and tidal changes. By using precise co-registration tools and change detection methods that exploit phase as well as amplitude levels of the backscattered signals, it was shown that coherent change detection can be utilized in this environment over time scales ranging from hours through several days.

Robustness of the fine-scale co-registration algorithm was demonstrated by testing its efficacy with image pairs that have been shifted and contaminated with noise. As applied to the data sets collected it was found that the SIFT algorithm was robustly able to co-register these contaminated image pairs, with the process breaking down only for high-noise exemplars. This breakdown was more pronounced for the feature-poor environment. Analysis of the distribution of the CCA coherence measure over datasets collected with different time latencies showed that coherency between image pairs that were co-registered with the fine-scale SIFT algorithm degraded over time; however, the sub-pixel co-registration algorithm restored some of the lost coherence – in many cases sufficiently to use the phase portion of the data in assessment of scene change.

Determination of seabed change, with levels of precision and robustness sufficient for alerting to the introduction of small man-made objects, is an extremely challenging problem. This paper demonstrated the potential of using automated image-based change detection techniques, made possible with the advent of synthetic aperture sonars and actively navigated undersea vehicles. The advantage of utilizing the phase component of these coherent sonars is that a detected change in phase between scenes indicates a change in bottom relief, even when there is negligible change in the amplitude of backscattered energy. This can be useful for investigating changes in topography or alerting to the presence of an acoustically chameleonic object. While the utility of exploiting changes in signal phase degrades over time, with time scales dependent on the environment, the capability of using image-based tools such as these to conduct incoherent change detection for time scales on the order of weeks has been demonstrated.[17]

A number of emerging tools exist to detect seabed changes in multi-temporal sidescan sonar images: automated image correlation – both coherent and incoherent, automated contact correlation, as well as interactive

tools to assist a sonar operator. For the near future performing seabed change detection over a variety of environments will likely require exploiting a number of these techniques – utilizing certain techniques for specific environments or combining a number of tools in parallel. The next steps in this research include: segmentation of detected changes to assess the nature and characteristics of detected regions of interest; reduction and rejection of clutter regions such as the water column and areas of decorrelation due to environmental processes; analytical assessment of which environments are most appropriate for this technique as well as effective combination with complementary techniques; and finally, the fast implementation of these image analysis algorithms on computer platforms.

7. ACKNOWLEDGMENTS

The authors gratefully acknowledge guidance and support from Dr. Jason Stack, Office of Naval Research Code 321OE, Ocean Engineering and Maritime Systems.

REFERENCES

[1] J. R. Anderson, "Land use and land cover changes, a framework for monitoring," *Journal of Research U.S. Geological Survey* **5**, pp. 143–153, March-April 1977.

[2] A. Singh, "Digital change detection techniques using remotely-sensed data," *Int. J. Remote Sensing* **10**(6), pp. 989–1003, 1989.

[3] E. J. M. Rignot and J. J. van Zyl, "Change detection techniques for ERS-1 SAR data," *IEEE Transactions on Geoscience and Remote Sensing* **31**, pp. 896–906, July 1993.

[4] A. Venot, J. L. Golmard, J. F. Lebruchec, L. Pronzato, E. Walter, G. Frij, and J. C. Roucayrol, "Digital methods for change detection in medical images," *Info. Processing in Medical Imaging, Proc. of the 8th Conference, 1983* , pp. 1–16, 1984.

[5] L. Lemieux, U. Wieshmann, N. Moran, D. Fish, and S. Shorvon, "The detection and significance of subtle changes in mixed-signal brain lesions by serial MRI scan matching and spatial normalization," *Med. Image Anal* **2**, pp. 227–242, September 1998.

[6] M. Bosc, F. Heitz, J. P. Armspach, I. Namer, D. Gounot, and L. Rumbach, "Automatic change detection in multimodal serial MRI: Application to multiple sclerosis lesion evolution," *NeuroImage* **20**, pp. 643–656, October 2003.

[7] D. Brie, M. Tomczak, H. Oehlmann, and A. Richard, "Gear crack detection by adaptive amplitude and phase demodulation," *Mechanical Systems and Signal Processing* **11**, pp. 149–167, January 1997.

[8] R. Bansal and D. Kazakos, "An algorithm for detecting a change in a stochastic process," *IEEE Trans. Info. Theory* **32**, pp. 227–235, March 1986.

[9] C.-Y. Fang, S.-W. Chen, and C.-S. Fuh, "Automatic change detection of driving environments in a vision-based driver assistance system," *IEEE Trans. Neural Networks* **14**, pp. 646–657, May 2003.

[10] D. Geronimo, A. M. Lopez, A. D. Sappa, and T. Graf, "Survey of pedestrian detection for advanced driver assistance systems," *IEEE Trans. Patt. Analy. and Machine Int.* **32**, pp. 1239–1258, May 2010.

[11] V. Myers, A. Fortin, and P. Simard, "An automated method for change detection in areas of high clutter density using sonar imagery," *Proc. Underwater Acoustic Measurements (UAM): Technologies and Results (UAM2009)* , 2009.

[12] P. Gamba, F. Dell'Acqua, and G. Lisini, "Change detection of multitemporal SAR data in urban areas combining feature-based techniques," *IEEE Transactions on Geoscience and Remote Sensing* **44**, pp. 2820–2827, Oct. 2006.

[13] M. Gendron, M. Lohrenz, G. Layne, and J. Ioup, "Automatic change detection and classification (ACDC) system," *Proc. of the Sixth International Symposium on Technology and the Mine Problem* , May 2004.

[14] D. D. Sternlicht, J. K. Harbaugh, and M. A. Nelson, "Experiments in coherent change detection for synthetic aperture sonar," *OCEANS 2009, MTS/IEEE Biloxi* , pp. 1–5, Oct. 2009.

[15] T. G-Michael and J. D. Tucker, "Canonical correlation analysis for coherent change detection in synthetic aperture sonar imagery," *Institute of Acoustics Proceedings* **32**, pp. 117–122, Sept. 2010.

[16] C. A. Matthews and D. Sternlicht, "Seabed change detection in challenging environments," *Proc. SPIE: Detection and Sensing of Mines, Explosive Objects, and Obscured Targets XVI* **8017**, pp. 80170P–80176P, May 2011.

[17] O. Midtgaard, R. E. Hansen, T. Saebo, V. Myers, J. Dubberley, and I. Quidu, "Change detection using synthetic aperture sonar: preliminary results from the larvik trial," *OCEANS 2011, MTS/IEEE* , pp. 1–8, September 2011.

[18] M. R. Azimi-Sadjadi and S. Srinivasan, "Coherent Change Detection and Classification in Synthetic Aperture Radar Imagery Using Canonical Correlation Analysis," *Internal Report, Information System Technologies, Inc.* , 2007.

[19] D. D. Sternlicht and T. G-Michael, "Change detection by image correlation for synthetic aperture sonar," *Institute of Acoustics Proceedings* **32**, pp. 87–91, Sept. 2010.

[20] J. D. Tucker and M. R. Azimi-Sadjadi, "Coherence-based underwater target detection from multiple sonar platforms," *IEEE J. of Ocean Eng.* **36**, pp. 37–51, Jan 2011.

[21] D. G. Lowe, "Object recognition from local scale-invariant features," *Proc. of the Seventh IEEE International Conference on Computer Vision* **2**, pp. 1150–1157, 1999.

[22] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision* **60**, pp. 91–110, Nov. 2004.

[23] D. D. Sternlicht, J. E. Fernandez, R. Holtzapple, D. P. Kucik, T. C. Montgomery, and C. M. Loeffler, "Advanced sonar technologies for autonomous mine countermeasures," *OCEANS 2011, MTS/IEEE* , pp. 1–5, September 2011.

[24] T. Saebo, R. E. Hansen, H. J. Callow, and S. A. V. Synnes, "Coregistration of synthetic aperture sonar images from repeated passes," *Proc. UAM: Technologies and Results (UAM2011)* , 2011.

[25] T. O. Saebo, *Seafloor depth estimation by means of interferometric synthetic aperture sonar*. Dissertation for phd, University of Tromso, Faculty of Science and Technology, Dept. of Physics and Technology, September 2010.

[26] P. Prats and J. Mallorqui, "Estimation of azimuth phase undulations with multisquint processing in airborne interferometric SAR images," *IEEE Trans. on Geoscience and Remote Sensing* **41**, pp. 1530–1533, June 2003.

[27] L. Scharf and C. Mullis, "Canonical coordinates and the geometry of inference, rate, and capacity," *IEEE Trans. Signal Process.* **48**, pp. 824–891, March 2000.

[28] A. Pezeshki, L. L. Scharf, J. K. Thomas, and B. D. Van Veen, "Canonical coordinates are the right coordinates for low-rank Gauss-Gauss detection and estimation," *IEEE Trans. Signal Process.* **54**, pp. 4817–4820, Dec 2006.

[29] A. P. Lyons and D. C. Brown, "The impact of the temporal variability of seafloor roughness on synthetic aperture sonar repeat-pass interferometry," *IEEE J. of Ocean Eng.* **38**, pp. 91–97, January 2013.