

Image-Based Automated Change Detection for Synthetic Aperture Sonar by Multistage Coregistration and Canonical Correlation Analysis

Tesfaye G-Michael, *Senior Member, IEEE*, Bradley Marchand, J. Derek Tucker, *Member, IEEE*, Timothy M. Marston, Daniel D. Sternlicht, *Member, IEEE*, and Mahmood R. Azimi-Sadjadi, *Senior Member, IEEE*

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 that relates and approximates pixel locations of reference and repeat-pass data sets; fine-scale coregistration using the scale-invariant feature transform (SIFT) algorithm to match features between overlapping data sets; local coregistration that improves phase coherence; and finally, change detection utilizing a canonical correlation analysis (CCA) algorithm to detect changes. The method was tested using data collected with a high-frequency SAS in a sandy shallow-water environment. Successful results of this multistage change detection method are presented here, and the robustness of the techniques that exploit phase and amplitude levels of the backscattered signals is discussed. 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. Robustness of the coregistration methods and analysis of scene coherence over time is characterized by analysis of repeat pass as well as synthetically modified data sets.

Index Terms—Automated change detection, canonical correlation analysis (CCA), coherent change detection, coregistration, synthetic aperture sonar (SAS).

I. 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 systems, change detection is the process by which regions of interest are identified through the comparison of current data with historical. Traditionally, there are two approaches to change detection: image correlation and contact correlation. For contact-based change detection, contact regions, flagged by a detector or an operator, are compared with those stored in a historical database. Through matching new contacts with historical, changes are identified. Coiras *et al.* [11] discussed an end-to-end contact-based change detection mechanism, where an automatic detector is followed by a rigid data association method to determine the location of previously undetected contacts. For image-based change detection, entire image portions are compared to identify regions of change. Myers *et al.* [12], [13] 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, resource management, and enhancement of automatic target recognition algorithms. Successful implementation of automated change detection should 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 [14]. In the sonar research community, contact correlation ACD approaches were developed first [15]. This was due to a number of factors, such as challenges in 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 (SASs) and actively navigated platforms, have image correlation methods been shown to be feasible [16]–[21].

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 compo-

Manuscript received December 20, 2013; revised July 03, 2014 and February 11, 2015; accepted July 31, 2015. This work was supported by the U.S. Office of Naval Research (ONR).

Associate Editor: T. Sæbø.

T. G-Michael, B. Marchand, and D. D. Sternlicht are with the Naval Surface Warfare Center Panama City Division, Panama City, FL 32407-7001 USA (e-mail: tesfaye.g-michael@navy.mil; bradley.marchand@navy.mil; daniel.sternlicht@navy.mil).

J. D. Tucker is with Sandia National Laboratories, Albuquerque, NM 87185 USA (e-mail: jdtuck@sandia.gov).

T. M. Marston is with the Applied Physics Laboratory, University of Washington, Seattle, WA 98105 USA (e-mail: marston@apl.washington.edu).

M. R. Azimi-Sadjadi is with the Electrical and Computer Engineering Department, Colorado State University, Fort Collins, CO 80523 USA (e-mail: azimi@engr.colostate.edu).

Digital Object Identifier 10.1109/JOE.2015.2465631

nents: historical data retrieval, reference and repeat-pass coregistration, and detection of changes.

A historical database comprising SAS images with their associated geographical locations provides the baseline or reference. Given a specific repeat-pass image, images from this database that correspond to the same geographical location within 30 m of each other are retrieved for use in the following coregistration and change detection steps. All image geographical information for both the database and repeat-pass images are collected from the SAS platform.

Image-based ACD for SAS is generally categorized as incoherent or coherent. Incoherent ACD identifies changes in the mean backscatter power of a scene, whereas coherent ACD identifies changes in both amplitude and phase of the transduced imagery (a process where the sound energy is converted into imagery through beamforming). In both cases, the detected changes are attributed to changes brought into the scene between data collection intervals. The phase component, present in coherent change detection (CCD), conveys more information regarding signal structure than magnitude alone. Oppenheim and Lim [22] demonstrated that the phase-only image retains many of the features as compared to the magnitude-only images in image synthesis. Furthermore, Huang *et al.* [23] presented the importance of phase in image processing filters. The advantage of utilizing the phase component is that a detected change indicates a physical disturbance in the scene, even when there is negligible change in the amplitude of backscattered energy. For SAR, it is generally accepted that processing the complex imagery is required for detecting very subtle man-made scene changes, such as vehicle tracks [24, pp. 19–20]. As conventional sidescan sonars exhibit resolution that degrades with range and are typically limited to the creation of backscatter amplitude images, they are less suitable for small-scale image-based ACD, thus explaining early adaption of contact-based ACD for sidescan systems.

Recent studies on temporal correlation of SAS data collected over various time intervals showed that temporal coherence holds up over time scales necessary for application of image-based change detection [19], [21]. Jackson *et al.* [25] described a model for the complex correlation between seafloor echoes acquired at different times (referred to as ping-to-ping correlation) based on estimating the roughness spectrum of sandy seafloors. Gerig *et al.* [26] studied the decorrelation of acoustic scattering as a function of temporal changes in seafloor roughness by comparing correlation estimates for pairs of data sets acquired from the same location, with identical instrument settings but at different times. In the context of change detection, Lyons and Brown [27] estimated the decorrelation of acoustic signals scattered from the seafloor due to changes in the shape of the water–sediment interface: providing insight into the impact of temporal variability of seafloor roughness on SAS repeat-pass images, and expressed temporal coherence as a function of acoustic wavelength.

The inspiration for the current work comes from an extension of the coherence-based change detection results using canonical correlation analysis (CCA) described by Azimi-Sadjadi and Srinivasan [28], G-Michael and Tucker [17], and Sternlicht and

G-Michael [29], 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 [30]. In this paper, we develop a complete ACD algorithm that uses the scale-invariant feature transform (SIFT) algorithm [31], [32] and local phase-based coregistration as our image registration tool followed by a CCA-based change detection method. In CCA, the correlation is formed between two sonar images from repeat passes over the ocean floor. From these correlations, coherence (or incoherence) can be measured and used to determine what has changed between the two sonar images at the same location but different times. Specifically, we define statistical relationship between two random variables and coherence to be a measure of dependence or strength of association between two signals.

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 (Panama City, FL, USA) and the Applied Research Laboratory, Pennsylvania State University (State College, PA, USA) [33]. The SSAM sonar is mounted on a Woods Hole Oceanographic Institution (WHOI, Woods Hole, MA, USA) 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. For fine-scale navigation and motion estimation, it is equipped with an onboard navigation system, which comprises a Kearfott KN-4902 Doppler velocity log and global positioning system (GPS)-aided inertial navigation system (INS), with a position error typically less than 0.2% of distance traveled. For multi-temporal data sets analyzed in this research, navigation offsets of less than 10 m were typical.

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 data sets; 2) fine-scale coregistration using the SIFT algorithm to match features between overlapping data sets; 3) local coregistration through optimizing the interscene phase coherence; and 4) coherent change detection utilizing the CCA algorithm to detect scene changes.

This paper is organized as follows. Section II provides detailed explanation of coarse and fine-scale coregistration method using SIFT. Section III describes the local coregistration step. Section IV describes the CCA algorithm and how it measures change. Section V presents a discussion of limitations and constraints of ACD method. Section VI presents performance results of our coregistration and change detection methods. Finally, Section VII provides conclusions, discussion, and directions for future research.

II. COREGISTRATION BY NAVIGATIONAL ALIGNMENT AND SIFT

Image coregistration is a key step in the ACD process. It provides a geometrical transformation that aligns the corresponding pixels from multitemporal images that represent the same area for comparison and change detection. For our coregistration method, we developed a three-stage process: navigation alignment, fine-scale coregistration, and local coregistration.

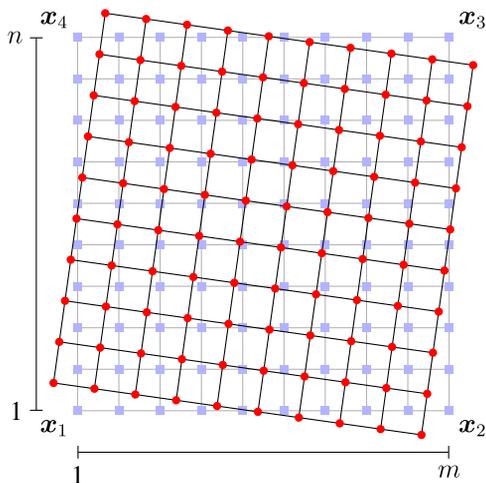


Fig. 1. Depiction of coarse alignment. Red circle and blue square points represent reference and repeat-pass data samples, respectively. Points \mathbf{x}_1 , \mathbf{x}_2 , \mathbf{x}_3 , and \mathbf{x}_4 are latitude and longitude position vectors for the repeat-pass image corners. Using the latitude/longitude to pixel mapping derived from (1) and (2), the reference data samples (red circles) are mapped into the repeat-pass pixel coordinate system. Now the repeat-pass data samples (blue squares) can be interpolated to the mapped reference locations (red circles) to recover the repeat-pass pixel values at reference data sample locations.

A. Navigation Alignment

After the identification and loading of the reference data corresponding to the repeat-pass image, a purely navigational coarse alignment is our first step toward coregistration 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 relative pixel location by solving the following system of equations:

$$\begin{bmatrix} 1 \\ m \\ m \end{bmatrix} = \begin{bmatrix} \mathbf{x}_1^\top & 1 \\ \mathbf{x}_2^\top & 1 \\ \mathbf{x}_3^\top & 1 \end{bmatrix} \mathbf{c}_1 \quad (1)$$

$$\begin{bmatrix} 1 \\ 1 \\ n \end{bmatrix} = \begin{bmatrix} \mathbf{x}_1^\top & 1 \\ \mathbf{x}_2^\top & 1 \\ \mathbf{x}_3^\top & 1 \end{bmatrix} \mathbf{c}_2 \quad (2)$$

where the size of the repeat-pass data is $m \times n$, $\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$ are the latitude and longitude position vectors of the repeat-pass image corners, and $\{\mathbf{c}_1, \mathbf{c}_2\}$ are the linear mapping coefficients that we are solving for. Using this mapping we can calculate the reference data pixel locations relative to the repeat-pass data, which allows interpolation of the repeat-pass data to the reference data pixel location. This method can result in rotation and translation of the repeat-pass image. The interpolation technique carried out in this paper is a standard linear interpolation. This projects the repeat-pass data on to the reference, thereby removing most navigation differences between reference and repeat-pass images. A depiction of this mapping is shown in Fig. 1.

B. Fine-Scale Coregistration

For fine-scale coregistration we apply the SIFT algorithm [31], [32] to the amplitude image data. SIFT geometrically

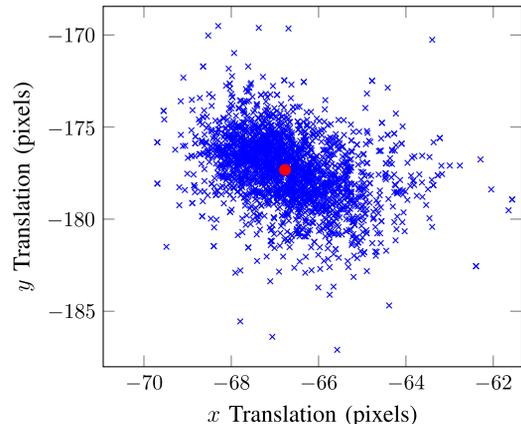


Fig. 2. Plot of translational shift estimates in pixels from SIFT matches (shown in blue). Red point indicates the determined shift estimated using the weighted average method described above. While shift estimates vary greatly due to incorrect SIFT matches, there is typically a significant portion of estimates with low variance, which enables good shift estimation.

aligns two images that represent the same scene but collected at different times. It is a feature-based technique consisting of feature detection and feature matching. The attractiveness of SIFT is mainly due to its invariance to various image transformations including: rotation, scaling, and displacements of pixels in a local region. The SIFT algorithm includes three steps: keypoint detection, descriptor establishing, and image feature matching.

For SIFT, image keypoints are interesting features or areas of an image, usually edges and corners in the image. A stable rotation and scale-invariant descriptor for each keypoint is created from the local image gradient. These descriptors are then used to carry out the fine-scale coregistration by matching candidate keypoints from reference and repeat-pass data (typically 10 000 keypoints for a SAS image). In this paper, we use the SIFT implementation by Vedaldi and Fulkerson [34].

Keypoint matching between reference and repeat-pass images is done by computing minimum Euclidean distance between SIFT features. To reject weak keypoints, matches where the second closest neighbor is within 1.5 times the distance to the closest match are rejected.

Differencing the matched keypoints typically results in a collection of noisy estimates for translational shifts. Fig. 2 shows an example from real SAS data being coregistered. There are many largely varying translation estimates, however there usually is a strong cluster which are much less varied. In Fig. 2, this cluster's center is indicated by the red point.

To identify this cluster and estimate the image translational shift, we first remove outliers by iteratively removing shift estimates three standard deviations away, then repeating. Once outliers have been removed, our final translation estimate is computed as a weighted average of the remaining translation estimates: specifically, letting $\{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_n\}$ be our remaining collection of shift estimates and c_{ij} be the Euclidean distance from \mathbf{s}_i to \mathbf{s}_j . We compute the weights $\{w_1, w_2, \dots, w_n\}$ as

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

For weights where $\min_{j \neq i} c_{ij}$ is zero (e.g., the shifts are the same), we set the weight value to twice the maximum of the defined weights to prevent an infinite cutoff value. The weights are then normalized

$$w'_i = \frac{w_i}{\sum_{j=1}^N w_j}$$

and the image shift is estimated as

$$\mathbf{s} = \sum_{i=1}^N w'_i \cdot \mathbf{s}_i. \quad (3)$$

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 environments that lack structured bottom features (e.g., smooth sandy bottom) and longer elapsed times between passes, the quality of the resulting estimate is not as good as the estimate obtained with the weighted average. Specifically, the estimate is more accurate when more structure is on the bottom to cross-reference the two images.

With the fine coregistration 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.

III. LOCAL COREGISTRATION THROUGH OPTIMIZATION OF INTERSCENE PHASE COHERENCE

Following fine-scale coregistration, the average scene displacement is ideally on the order of a single pixel. Local registration error, however, can be significantly higher than the average. This is particularly true if the trajectories used to form the reference and matching images are significantly different. A large difference in altitude relative to the ground plane between repeat passes, for example, results in an interferometric baseline that has associated with it all the same range-variant footprint shift effects described in the interferometric literature [35] that cause decorrelation. Some local translation or phase differences, on the other hand, may actually be caused by environmental features or changes. The retention of this information is critical. It is, therefore, important to be able to distinguish between local coregistration errors that are caused by environmental features or changes versus those that are caused by uncompensated navigation differences.

In this section, a method is outlined by which the local coregistration errors are estimated and fit to a set of functions that have simple physical interpretation as navigation errors. The residual difference between the estimated model and the data is dominated by fine-scale environment and change information. In this manner, navigation errors are separated from environmental features and compensated for, increasing correlation by refining local coregistration and isolating the relevant

repeat-pass information necessary for change detection. An outline of the algorithm steps is shown in Algorithm 1 and the steps are elaborated on in Sections III-A–III-G.

This approach, in contrast to that described in [36], operates only on the beamformed data rather than raw data and the process is noniterative in the sense that the 2-D interpolation and phase correction are only applied once. As has been observed, autofocusing the beamformed data to correct for focusing errors has a positive impact on the correlation between reference and matching images. An additional step of autofocusing was applied to all the examples in this paper. The autofocus algorithm was the patch-wise application of multiple-aperture map-drift [37, Ch. 6], which was modified for wideband, widebeam systems. Patches were overlapped and the linear component of the phase solutions was removed to prevent shifting of features and discontinuities between patch boundaries. A demonstration of how autofocusing affects the local coherence is provided in Section VI.

A detailed step-by-step description of the local coregistration algorithm is presented as follows.

Algorithm 1: Local coregistration

Step 1. Large neighborhood (e.g., 50×50 pixel) nonoverlapping patch correlation is performed.

Step 2. Correlation peaks for each patch are found and parabolic interpolation is used to estimate the local relative patch translations in the along-track (x) and across-track (y) dimensions to local precision.

Step 3. Local along-track translation measurements are parametrized as a series of coarse surge and heading errors relating the geometries of the synthetic apertures. Across-track translation measurements are parametrized as a series of heave (vertical translation) and sway (horizontal translation) errors.

Step 4. The surge, heading, sway, and heave vectors are used to form 2-D interpolation functions for remapping the matching image onto the same grid as the reference image, correcting for local registration errors.

Step 5. The complex correlation coefficient between images is calculated using a sliding, small neighborhood (e.g., 10×10 pixels). The phase of the complex correlation forms an interferogram which is unwrapped.

Step 6. The interferogram is decomposed into surge and sway functions.

Step 7. The surge and sway functions are used to define the argument of a phase function that, when multiplied by the matching image, removes the effects of surge and sway on the interferogram formed between the reference and matching images.

A. Large Neighborhood Correlation

The first step of the local coregistration process is to estimate the local coregistration errors in terms of their along-track (x -dimension) and across-track (y -dimension) components. This is done by dividing the reference image, denoted as complex 2-D

matrix $A(x, y)$, and matching image $B(x, y)$ into $I \times J$ nonoverlapping 50×50 pixel neighborhoods which are correlated using 2-D Fourier transforms, e.g.,

$$Z_{i,j} = \frac{\mathcal{F}^{-1}(\mathcal{F}(A_{i,j}) \circ \mathcal{F}(B_{i,j})^*)}{\sqrt{(\sum |A_{i,j}|^2)(\sum |B_{i,j}|^2)}}. \quad (4)$$

In (4), the indices i and j range from 1 to I and 1 to J and \circ is the Hadamard product operator. The 2-D discrete Fourier transform is denoted by \mathcal{F} , and $(\cdot)^*$ denotes the complex conjugate. The matrix $Z_{i,j}$ contains a peak that has a magnitude ranging from zero to one, with greater values implying stronger correlation. The location of the peak indicates the relative translational offset between patches $A_{i,j}$ and $B_{i,j}$.

B. Subpixel Local Translation Estimation

Translation measurements in x and y are found to subpixel precision by fitting a 1-D parabola to the three points formed by the peak value in $Z_{i,j}$ and the adjacent neighboring pixels. The location in x or y of the parabola peak is chosen as the refined translational offset. If $Z_{i,j}(x_0, y_0)$ is the peak value of $Z_{i,j}$, then the x and y translations are estimated as

$$\begin{aligned} T_x(i, j) &= x_0 - \frac{1}{2} (Z_{i,j}(x_0 + 1, y_0) - Z_{i,j}(x_0 - 1, y_0)) \\ &\quad \times \frac{1}{Z_{i,j}(x_0 + 1, y_0) - 2 \cdot Z_{i,j}(x_0, y_0) + Z_{i,j}(x_0 - 1, y_0)} \quad (5) \\ T_y(i, j) &= y_0 - \frac{1}{2} (Z_{i,j}(x_0, y_0 + 1) - Z_{i,j}(x_0, y_0 - 1)) \\ &\quad \times \frac{1}{Z_{i,j}(x_0, y_0 + 1) - 2 \cdot Z_{i,j}(x_0, y_0) + Z_{i,j}(x_0, y_0 - 1)}. \quad (6) \end{aligned}$$

If the scenes have intrinsically low coherence, demonstrated by a low average peak values for $Z_{i,j}$ over all I and J , the large neighborhood correlation and subpixel translation estimates are recomputed using only the magnitude of A and B . In the present implementation, this alternate method is attempted if the average correlation coefficient is < 0.2 .

C. Parametrization of Translation Estimates

Large neighborhood correlation results in three matrices: a matrix of along-track (x) translations (T_x), range (y) translations (T_y), and a map of correlation coefficient values ($|Z_{i,j}|$) indicating the quality of the translation estimates. The next steps fit the translation matrices using a set of functions designed to model navigation errors. The goal of this fitting is to find a 2-D interpolation function that compensates for the local displacement errors of the matching image due to the effects of the sonar platform having traveled along a different path during synthetic aperture formation. Constraining the correction to be modeled as a set of navigation errors prevents overfitting, which could potentially remove important environmental change information.

- Along-track translation parametrization: Along-track errors are modeled as having two components: surge and heading (i.e., rotation). Surge error is defined as synthetic

aperture sampling error in the along-track direction. In the current implementation, its effects are approximated as constant in the range dimension. It is, however, unconstrained in the along-track dimension. Rotation errors are assumed to be small and are modeled as an along-track translation that varies linearly with range. Combining these two, the full model for along-track translation is a first-order polynomial plus noise

$$T_x(i, j) = r(j)\mathbf{t}_{\text{rot}}(i) + \mathbf{t}_{\text{surge}}(i) + N(i, j) \quad (7)$$

where i is the patch index in the along-track dimension, j is the patch index in the range dimension, \mathbf{t}_{rot} is a vector of the along-track translation values in pixels due to rotation, $\mathbf{t}_{\text{surge}}$ is a vector of along-track translations due to surge, $r(j)$ is range of patch j in pixels, and $N(i, j)$ is a noise matrix.

The level of noise can be estimated via the correlation coefficient values of $Z_{i,j}$, and it is rarely uniform. To fit the model to the measured translation estimates, an iterative weighted least squares approach is utilized. We set our rotation fitting function $\boldsymbol{\zeta}_{\text{rot}}$ to be a linearly varying vector from zero to one and our surge fitting function $\boldsymbol{\zeta}_{\text{surge}}$ to be a vector of ones. Both are of length J . The weighting function is then initialized with the values from the correlation coefficient matrix, and a weighted pseudoinverse is used to fit the functions, e.g.,

$$\begin{bmatrix} \hat{\mathbf{t}}_{\text{rot}}(i) \\ \hat{\mathbf{t}}_{\text{surge}}(i) \end{bmatrix} = \left(\begin{bmatrix} \boldsymbol{\zeta}_{\text{rot}}^T \\ \boldsymbol{\zeta}_{\text{surge}}^T \end{bmatrix} \left[(\mathbf{w} \circ \boldsymbol{\zeta}_{\text{rot}}) (\mathbf{w} \circ \boldsymbol{\zeta}_{\text{surge}}) \right] \right)^{-1} \times \begin{bmatrix} (\mathbf{w} \circ \boldsymbol{\zeta}_{\text{rot}})^T \\ (\mathbf{w} \circ \boldsymbol{\zeta}_{\text{surge}})^T \end{bmatrix} T_x(i)^T \quad (8)$$

where $\mathbf{w} = |Z_{i,j}(i)|^T$ is a weighting vector defined as the i th row of correlation values from the correlation coefficient matrix and $T_x(i)$ is the i th row of the along-track translation measurement matrix. Though the results found by applying (8) represent an improvement over the unweighted version, the correlation coefficient values often overpredict the quality of a particular translation estimate, especially when the quality is poor. For this reason, outliers can still have a strong impact on the result. To reduce the effects of outliers the weights are iteratively redefined as

$$\mathbf{w}_k = \frac{\mathbf{w}}{\left(\hat{\mathbf{t}}_{k-1} - T_x(i) \right)^2 + \sigma} \quad (9)$$

where k is the current iteration number, σ is a constant that controls the level of outlier rejection ($\sigma = 1$ in the current examples), and $\hat{\mathbf{t}}_{k-1}$ is the signal reconstructed from the surge and rotation translation estimates found from the previous iteration, e.g.,

$$\hat{\mathbf{t}}_{k-1} = \hat{\mathbf{t}}_{\text{rot}}(i) \circ \boldsymbol{\zeta}_{\text{rot}}^T + \hat{\mathbf{t}}_{\text{surge}}(i) \circ \boldsymbol{\zeta}_{\text{surge}}^T. \quad (10)$$

- Across-track parametrization: In prior literature, it has been noted that the range variance of the temporal displacement between reference and matching signals varies approximately proportional to the cosine of the grazing

angle in the case of horizontal offsets and the sine of the grazing angle for vertical offsets [38]. Thus, the range displacement can be coarsely defined as the linear combination of sway (horizontal) and heave (vertical) components and noise, e.g.,

$$T_y(i, j) = \mathbf{t}_{\text{heave}}(i) \sin(\theta_g) + \mathbf{t}_{\text{sway}}(i) \cos(\theta_g) + N(i, j) \quad (11)$$

where $T_y(i, j)$ is a matrix of local range translation measurements in pixels, θ_g is the grazing angle defined here as the arctangent of the average synthetic aperture height $h(i)$ divided by the ground-plane range defined by making a flat-bottom assumption

$$\theta_g = \tan^{-1} \left(\frac{h(i)}{\text{Real} \left(\sqrt{r(j)^2 - h(i)^2} \right)} \right). \quad (12)$$

The estimates of $\hat{\mathbf{t}}_{\text{heave}}$ and $\hat{\mathbf{t}}_{\text{sway}}$ are found using the same weighted pseudoinverse and iterative reweighting scheme with the exception that the fitting functions are redefined as ζ_{heave} and ζ_{sway} , the sine and cosine of the grazing angle.

D. Image Remapping

The $\mathbf{t}_{\text{surge}}$, \mathbf{t}_{rot} , $\mathbf{t}_{\text{heave}}$, and \mathbf{t}_{sway} vectors containing the estimated navigation errors can be used in conjunction with their corresponding function expressions $\zeta_{\text{surge}}(r)$, $\zeta_{\text{rot}}(r)$, $\zeta_{\text{heave}}(r)$, and $\zeta_{\text{sway}}(r)$ to construct a 2-D surface defining an interpolation map for correcting local registration errors in the matching image. Specifically

$$\hat{T}_x(i, j) = \mathbf{t}_{\text{rot}}(i) \circ \zeta_{\text{rot}}^T(r(j)) + \mathbf{t}_{\text{surge}}(i) \circ \zeta_{\text{surge}}^T(r(j)) \quad (13)$$

$$\hat{T}_y(i, j) = \mathbf{t}_{\text{sway}}(i) \circ \zeta_{\text{sway}}^T(r(j)) + \mathbf{t}_{\text{heave}}(i) \circ \zeta_{\text{heave}}^T(r(j)) \quad (14)$$

form our translational shifts, which provides the following warping to our along-track (\mathbf{p}) and across-track (\mathbf{q}) pixel values:

$$\hat{\mathbf{p}} = \mathbf{p} + \hat{T}_x(j) \quad (15)$$

$$\hat{\mathbf{q}} = \mathbf{q} + \hat{T}_y(i)^T \quad (16)$$

where $\hat{T}_x(j)$ is the j th column of \hat{T}_x and $\hat{T}_y(i)$ is the i th row of \hat{T}_y .

The vectors are only of length I , however, and must be first upsampled to the length of the original complex image matrices. In the current implementation, this is done via linear interpolation. The interpolated vectors are then used to define 2-D along-track and across-track interpolation functions that remap the complex matching image to align with the reference image and correct for coregistration errors due to navigation discrepancies between synthetic apertures.

E. Small Neighborhood Correlation Interferogram

At this point in the local coregistration algorithm the features of the synthetic aperture images have been aligned. The phase, however, has not yet been modified to compensate for the navigation discrepancies. The next steps utilize the interferogram to

generate a phase correction for removal of the trends introduced by navigation discrepancies.

Following remapping, a 10×10 pixel sliding window is used to calculate complex correlation values between corresponding image neighborhoods

$$Z_{i,j} = \frac{\sum (A_{i,j} \circ B_{i,j})^H}{\sum \sqrt{|A_{i,j}|^2 \circ |B_{i,j}|^2}}. \quad (17)$$

The selection of the sliding window size is somewhat arbitrary; this choice over many data sets seemed to provide a good balance between speed of calculation, noise reduction, and preservation of detail before unwrapping. The phase of $Z_{i,j}$ forms an interferogram between the images. A variety of methods can be used to unwrap interferograms [39]. In this paper, an iterative weighted least squares approach is implemented for unwrapping. The reader is referred to [39] for more information on weighted least squares phase unwrapping. Because of the computational burden associated with weighted least squares phase unwrapping, the phase was smoothed and decimated by a factor of ten in both the x - and y -dimensions before unwrapping.

F. Heave and Sway Estimation

Heave and sway are reestimated using the unwrapped interferogram. Coarse estimates of heave and sway were estimated during Step 3 using the large neighborhood translation measurements, but a larger number of samples are available for fitting when using the unwrapped interferogram. The same iterative weighted least squares approach utilized in Step 3 is again used to decompose the unwrapped phase surface into sway and heave components.

Because the phase unwrapping method does not result in a solution that is correct on an absolute scale (i.e., the phase may have a static offset) the offset needs to be estimated and removed. To estimate this offset, an additional function (a static offset with range) is introduced during the first iteration. After fitting, the median of the static offset values is then subtracted from the phase surface and it is refitted using only the cosine and sine of the grazing angle.

G. Phase Correction

The resulting 2-D phase function described by heave and sway can be subtracted from the interferogram previously produced by correlation. The uncompensated phase trends reflecting navigation discrepancies between the matching and reference synthetic apertures can be compensated. The correction $\xi(\hat{\mathbf{p}}, \hat{\mathbf{q}})$ is calculated the same way as (14), except using our refined heave and sway estimates from the previous section. Our phase corrected image is then

$$\hat{B}(\mathbf{p}, \mathbf{q}) \mapsto B(\hat{\mathbf{p}}, \hat{\mathbf{q}}) e^{-j\xi(\hat{\mathbf{p}}, \hat{\mathbf{q}})}. \quad (18)$$

IV. CHANGE DETECTION

For change detection, we utilize the CCA [17] algorithm, which operates on complex SAS image data. CCA is a multivariate statistical method [40] 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 interchannel coherence. In other words, CCA finds a linear combination of the input vectors that have maximum correlation with each other.

Let us assume that the two-channel complex data vectors are represented by $\mathbf{x} \in \mathbb{C}^n$ and $\mathbf{y} \in \mathbb{C}^m$, where \mathbf{x} is the reference SAS image and \mathbf{y} is the repeat-pass image. The canonical correlations and coordinates are then obtained by finding the singular values and vectors of the coherence matrix [40], which is defined as

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

where $(\cdot)^H$ is the Hermitian operation, $E[\cdot]$ is the expectation operation, and R_{xx} and R_{yy} are the covariance matrices of \mathbf{x} and \mathbf{y} channels, respectively (in other words, \mathbf{x} and \mathbf{y} are random vectors with zero means and they yield a composite covariance matrix R_{xx} and R_{yy}). We block out the image into $n \times n$ blocks and reshape those blocks into a vector and concatenate into a data matrix $X = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n] \in \mathbb{R}^{n^2 \times N}$, where \mathbf{x}_i is a vector of the i th block. We then form the covariance using $R_{xx} = 1/N(X - \bar{X})(X - \bar{X})^T$, where \bar{X} is the mean vector of X .

As illustrated in the Appendix, the canonical coordinates of \mathbf{x} and \mathbf{y} are defined as

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

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 \mathbf{u} and \mathbf{v} share the diagonal cross-covariance matrix

$$K = R_{uv} = F^H C G \quad (19)$$

known as the canonical correlation matrix of canonical correlations $k_i, i \in [1, m]$, which are the singular values of matrix C . In other words, canonical correlation is a method for investigating the relationship between two groups of variables, by finding linear functions of one of the sets of variables that maximally correlate with linear functions of the variables in the set.

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

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

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

$$H = 1 - L, \quad 0 \leq H \leq 1. \quad (20)$$

That is, the channels \mathbf{x} and \mathbf{y} are perfectly coherent if $H = 1$, and noncoherent if $H = 0$. The reader is referred to the Appendix for a more detailed mathematical and theoretical formulation of CCA.

Though the dominant canonical coordinates capture maximal coherence, the subdominant ones maximize the divergence between the two channels \mathbf{x} and \mathbf{y} [17], [28], [41]. 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 \mathbf{v} and \mathbf{u} can be calculated using the residual, $\mathbf{v} - K\mathbf{u}$, owing to the fact that $K\mathbf{u}$ brings the most coherent information that the second channel \mathbf{y} carries about the first channel \mathbf{x} . Now that the subdominant canonical coordinates capture the coherent change information (under the constraint of maximum rate; the rate here refers to the information rate, the rate at which the canonical coordinate of \mathbf{x} channel carries information about the canonical coordinate of \mathbf{y} channel and *vice versa*), retaining the last few components of the difference vector $\mathbf{v} - K\mathbf{u}$ is sufficient. This can be performed mathematically by using $I_p(\mathbf{v} - K\mathbf{u})$, 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

$$\mathbf{d} = R_{yy}^{\frac{1}{2}} G I_p (\mathbf{v} - K\mathbf{u}). \quad (21)$$

For a more complete description of the CCA change detection method for SAS imagery, the reader is referred to [17].

V. LIMITATIONS AND CONSTRAINTS

The relevant repeat-pass time scales required for maintaining coherence are necessarily going to be highly dependent on the temporal variability of the environment in question. Some marine environments are inherently more variable due to the nature of the sediment (e.g., sand versus rock), while some experience high levels of biologic activity which cause scene decorrelation due to both motility and physical perturbation of the substrate (e.g., the artificial reef effect [27]). Some environments exhibit strong temporally variant currents which can rapidly deposit new sediment or change existing sediment structures. Preliminary studies have been conducted that measure and model correlation decay as a function of time and acoustic wavelength [27]; however, expanding these models to a broad variety of mechanisms and sedimentary environments is still an active area of research. The method for automated coherent change detection presented herein may be a useful tool for performing additional studies in a broad variety of environments, and without the artificial reef effect associated with the usage of more permanent structures.

With regard to the ACD method presented in this paper, the individual steps of our method have various limitations. To achieve an initial coarse alignment, in the multistage coregistration, we used nominal paths and linear mapping from latitude and longitude of reference and repeat-pass images. This step requires that the vehicle navigation offsets between passes be smaller than the size of the image patch being compared, otherwise no common information will be available

for coregistration. Further, the use of linear interpolation may introduce artificial misalignment. Fine coregistration, SIFT, is highly dependent on stable features that are common to both reference and repeat-pass images. The highly temporal nature of many of the features in a SAS image coupled with biological activity present a significant challenge for a feature-based coregistration method. Further, SAS images are highly dependent on aspect to the insonified area, which further imposes accurate navigation to maintain stable features. The local coregistration is contingent on there being some sort of mutual complex backscattering information present between the overlapping patches and insonifying the scene over the same wave number space. The primary constraints here are the multipath interference that adds noise to SAS image and obscure target structure.

VI. DATA AND RESULTS

To demonstrate the feasibility of our ACD approach, we operated on SAS data collected with the high-frequency projector of the SSAM system for a variety of measurement scenarios. The data collection procedure consisted of first running a linear sonar track to create the reference image, placing a target on the bottom, then repeating the originally programmed vehicle track to create the repeat-pass image. This was followed by additional runs at various azimuthal offsets to discern the sensitivity of the coregistration and change detection algorithms to geometric and radiometric distortion. The targets were deployed at approximately 25-m water depth off the coast of Panama City, FL, USA, during a period of sea state 2 on the Beaufort scale. We start with an analysis of the coregistration performance, followed by an analysis of the change detection performance.

A. Coregistration

Results using the initial coregistration steps described in Section II are shown in Figs. 3 and 4. In Fig. 3, the vehicle followed the same track with approximately 2 h between passes. To introduce change, a 208-L drum was dropped into the field between passes. Fig. 3 presents a relatively benign coregistration environment due to the short time period between passes and the high quantity of prominent structure available in the images. In these conditions, the initial steps of coarse alignment and SIFT are able to achieve high-precision coregistration of the SAS images. A more difficult coregistration environment can be seen in Fig. 4: a flat sandy region with sparse structure for image registration. Such regions are more difficult due to the temporal instability of the base image structure. In this case, with a short time latency of 2 h between passes, successful coregistration was possible. We were able to achieve high-precision automated coregistration in this sandy environment for time latencies of up to 72 h, then with more difficulty for longer time periods as is shown in the next section (Fig. 18).

To test the sensitivity of the coregistration process to variations in vehicle trajectory, data were collected from the same vicinity and time latency as in Fig. 3; however, in this case, the vehicle path of the repeat pass was 20° offset from that of the reference. Fig. 5 shows that, despite this path difference and the slight change in structure that results, the algorithm registers the images precisely. A navigational path that differs by 90° or

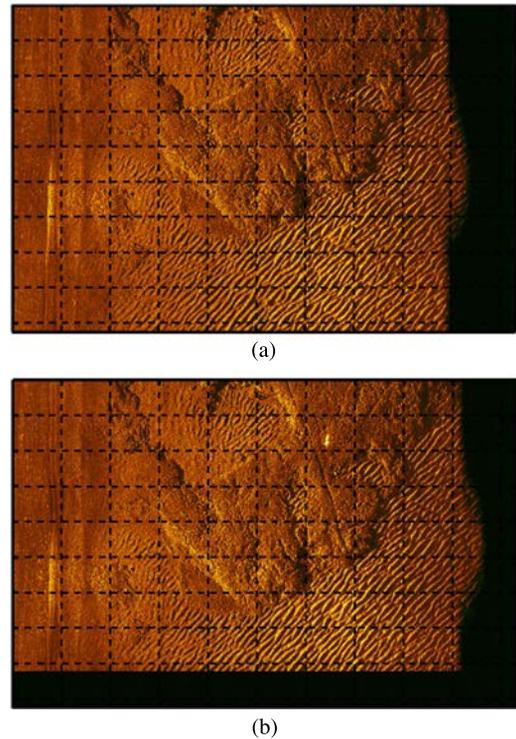


Fig. 3. Automated coregistration for data collected from data set 1. The vehicle was programmed to follow the same navigational path for both passes. A new object is present in the top right quadrant of the repeat-pass image. (a) Reference. (b) Repeat pass.

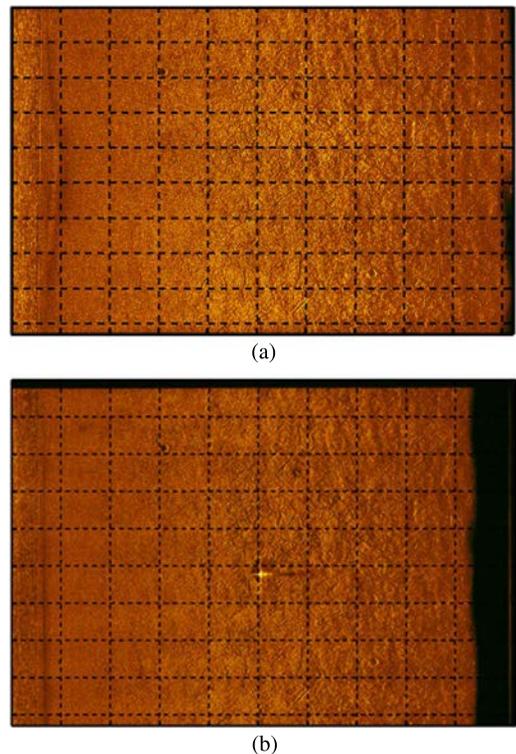


Fig. 4. Automated coregistration results collected from data set 2. A new object is present in the center of the repeat-pass image. (a) Reference. (b) Repeat pass.

180° from the reference path will result in images with substantially different structure, making image-based coregistration extremely difficult in the absence of angle-independent markers (see Fig. 6). This is because image coregistration requires that

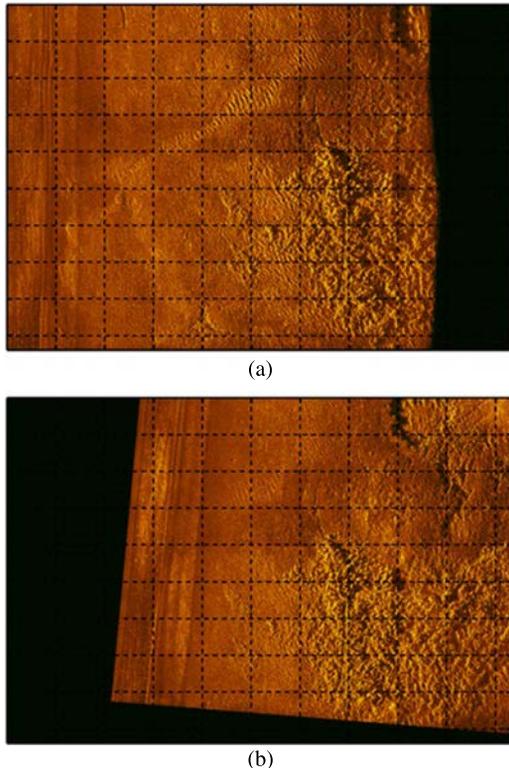


Fig. 5. Automated coregistration of data collected from data set 1. The vehicle was programmed to follow a navigational path that was 20° offtrack from the track used in Fig. 3. No new object is present in this scenario. (a) Reference. (b) Repeat pass.

spectral overlap exist in wave number space between the different surveys. If the path orientation is such that with the given system beamwidth none of the same spatial wave numbers for the reflectivity map are being populated, then coregistration and change detection are, by default, impossible.

The bandwidth, beamwidth, and beam-orientation of a synthetic aperture system all determine the bounded region of wave number space that the complex reflectivity function of the seafloor is sampled [42]. Realistic environmental conditions can change the orientation of the sonar platform and beam angle relative to its direction of motion. This effect, called crabbing, rotates the sampled wave number region around the origin of the spectrum [42]. If the sonar has a different crab angle during a repeat-pass scan, the associated wave number spectra for the reflectivity functions from both passes may have reduced overlap, causing decorrelation.

To provide some robustness to crabbing, in beamforming, the complex wave number spectrum can be filtered to accept data corresponding to signals arriving at angles plus or minus one quarter the full system beamwidth relative to the average heading of the synthetic array. This filtering allows for repeat-pass crab angle variability of plus or minus one quarter the full beamwidth relative to the average heading without loss of correlation. The drawback is that the resolution of the system is cut in half. Furthermore, crab angles beyond these limits still cause decorrelation.

1) *Local Coregistration Example:* Next, we will illustrate stepwise effects of the local coregistration process applied to

Fig. 3 to clarify the impact on the data. We begin with the along-track translation estimation of Steps 2 and 3 of Algorithm 1. The effects of subpixel refinement (Step 2) are shown in Fig. 7. In Fig. 7(a), the along-track correlation peak translation estimates are shown. Fig. 7(b) shows the subpixel refined estimate resulting from the application of (5). With the refined estimates we see much smoother and more likely initial estimates of translation. We can also see small regions where translation estimates are in disagreement with neighboring trends. These are likely regions of low correlation or image change between reference and repeat passes. These are changes we want to maintain while removing effects of navigational discrepancy, which is the goal of Step 3 shown in Fig. 8.

The resulting improvement in coherence from applying the model fitted translation estimate is shown in Fig. 9. The figure shows the small neighborhood correlation of the reference and repeat-pass images (Step 4) before and after remapping. Areas corresponding to sediment backscatter have high coherence, (0.9) throughout most of the image, the exceptions being in regions of low signal-to-noise ratio (SNR) such as shadows (caused by sand ripples), before bottom-detect (low range bins) and shallow grazing angles (high range bins).

Next, we correct errors in the phase using the same process as before at a finer scale using the small neighborhood interferogram generated from the previous step. The interferogram phase is first unwrapped, shown in Fig. 10. The interferogram [Fig. 10(a)] reveals that the navigation error contributes to phase noise and that the phase is dependent on the correct navigation alignment. This is exhibited by the colored bands or interferometric fringes, whose presence is an indication of interscene phase decorrelation. Our goal of applying the phase correction is to remove the quickly varying fringe patterns caused by navigation errors and obtain the result shown in Fig. 11(b).

The heave and sway are reestimated using the unwrapped interferogram (Step 6). Plots of the reestimated heave and sway and the detrended interferogram are shown in Fig. 11. The resulting phase differences are a result of bathymetry and possibly other effects such as local turbulence or temperature changes. The resulting 2-D phase function described by heave and sway can be subtracted from the interferogram previously produced by correlation shown in Fig. 11.

The results of the final phase correction (Step 7) is presented in Fig. 12. Noticeable are the lack of surge and rotation estimates. The wave number distribution in the along-track direction is centered around $k_x = 0$, whereas in the range direction the wave number k_y is centered around $k_y = 2\omega_0/c$, where ω_0 is 2π multiplied by the center frequency of the transmitted pulse in hertz. The fact that the along-track wave number spectrum straddles $k_x = 0$ results in a low level of phase sensitivity in the coherence function to small translations in the along-track direction. The result is that the interferogram is only sensitive to translations in the range dimension, namely those caused by sway and heave errors as displayed.

Last, Fig. 13 depicts the computed correlation coefficients of Fig. 4 with and without the application of the autofocus algorithm. We see that compensation of local focusing aberrations raises the coherence between images without artificially

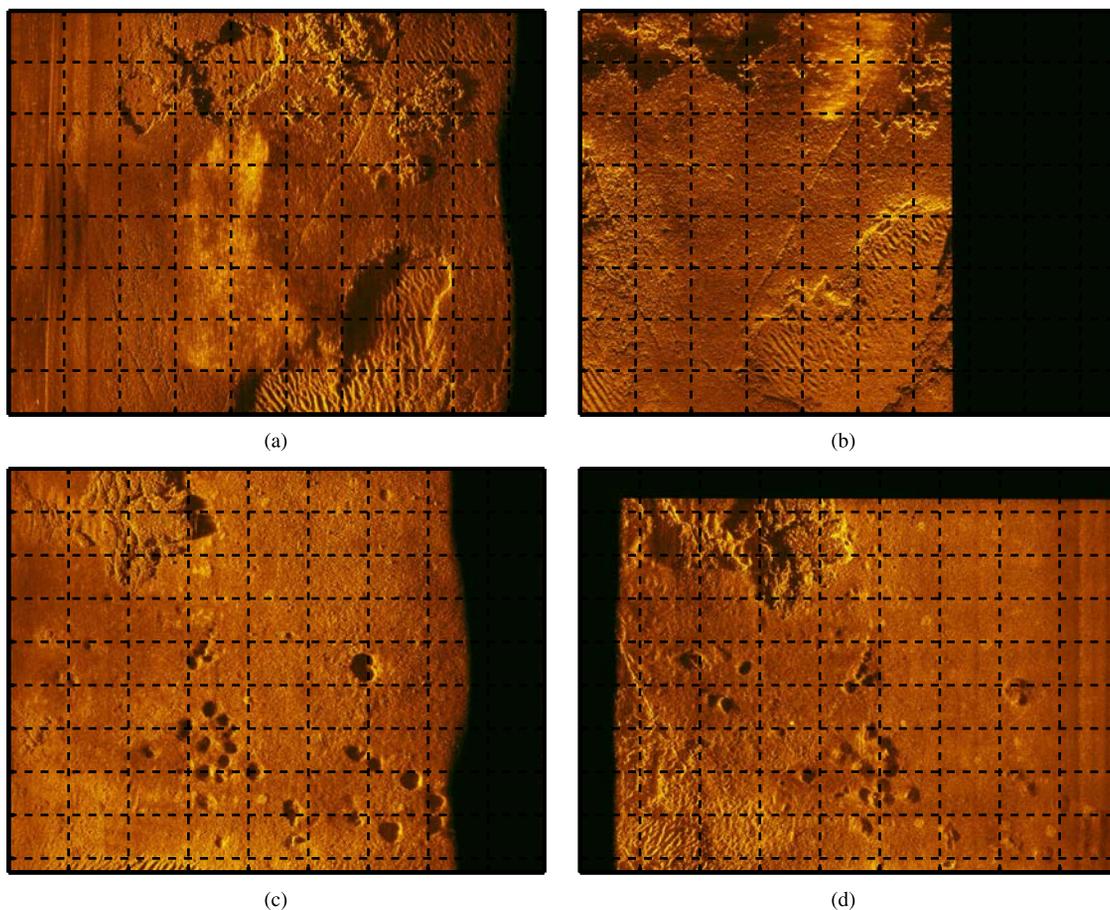


Fig. 6. Reference and repeat-pass images that are (a) and (b) 90° and (c) and (d) 180° offtrack from each other after coregistration. Local coregistration has not been applied.

raising the coherence of regions in the image that are known to be low, e.g., the large low-coherence region near the center corresponding to an object introduced into the scene still maintains low coherence.

2) *Robustness of the Coregistration Algorithm of Section II:* To quantify the robustness of our coregistration 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 coregistration algorithm can then be calculated as the Euclidean distance between the known image shift and the shift estimated by coregistration. 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. The results of applying the SIFT coregistration algorithm are shown in Table I. The table contains the mean error and standard deviation for noise levels and shift amounts considered. In general, the size of the shift difference does not impact coregistration performance for the shifts considered. For low and moderate noise, coregistration works well with typically less than a pixel of distance error in coregistration. For high noise, coregistration performs well for some images. When it does fail, the estimated shift is vastly different from the true shift, which significantly impacts the mean error as evident from the large standard deviations. To better illustrate this, Fig. 14 shows the percentage of images with estimated shift greater than the true shift values for the high noise

scenario of Table I. Almost half of all the errors are less than 20 pixels in size, while most of the other errors are greater than 200 pixels in size. The extremely large coregistration failures can likely be detected and flagged for operator intervention. Fig. 14 also shows the relative difficulty posed by the two environments. For the all-sand environment (data set 2), twice as many large coregistration failures occur at high noise, compared to the feature rich environment (data set 1).

B. Change Detection by CCA

To demonstrate the performance of the CCA change detection algorithm described in Section IV, we first processed the data sets through the multistage coregistration algorithms outlined in Sections II and III. The results of inputting two of the coregistered image pairs into the CCA modules are shown in Figs. 15 and 16. For the flat sandy environment of Fig. 4, the new object is clearly identified by the low-level region of the coherence map created with the technique of CCA (Fig. 16). For this example, CCD is not necessary to detect the object. The object can easily be detected with an incoherent change detection (ICD) method. One benefit of CCD is increasingly sensitive. In Fig. 17, we show an example that compares the coherent and incoherent (difference of log ratio) change detections between image patches where a change can be seen in the CCD, but not in the ICD. This example reveals clear distinction between CCD and ICD. The ICD is sensitive to significant changes, i.e., change which intensely influences the backscatterer of a target,

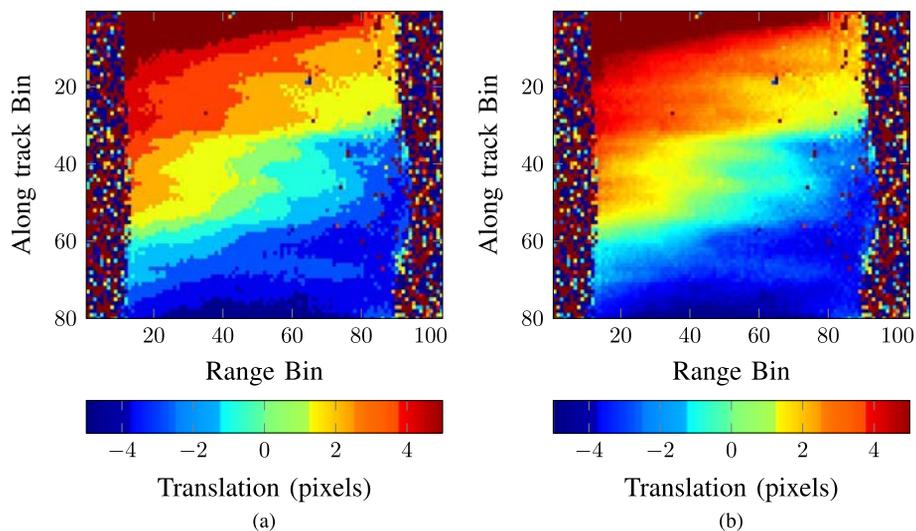


Fig. 7. Along-track translation measured in pixels. Image (a) utilizes the index values of the correlation peak, and image (b) is the refined estimate.

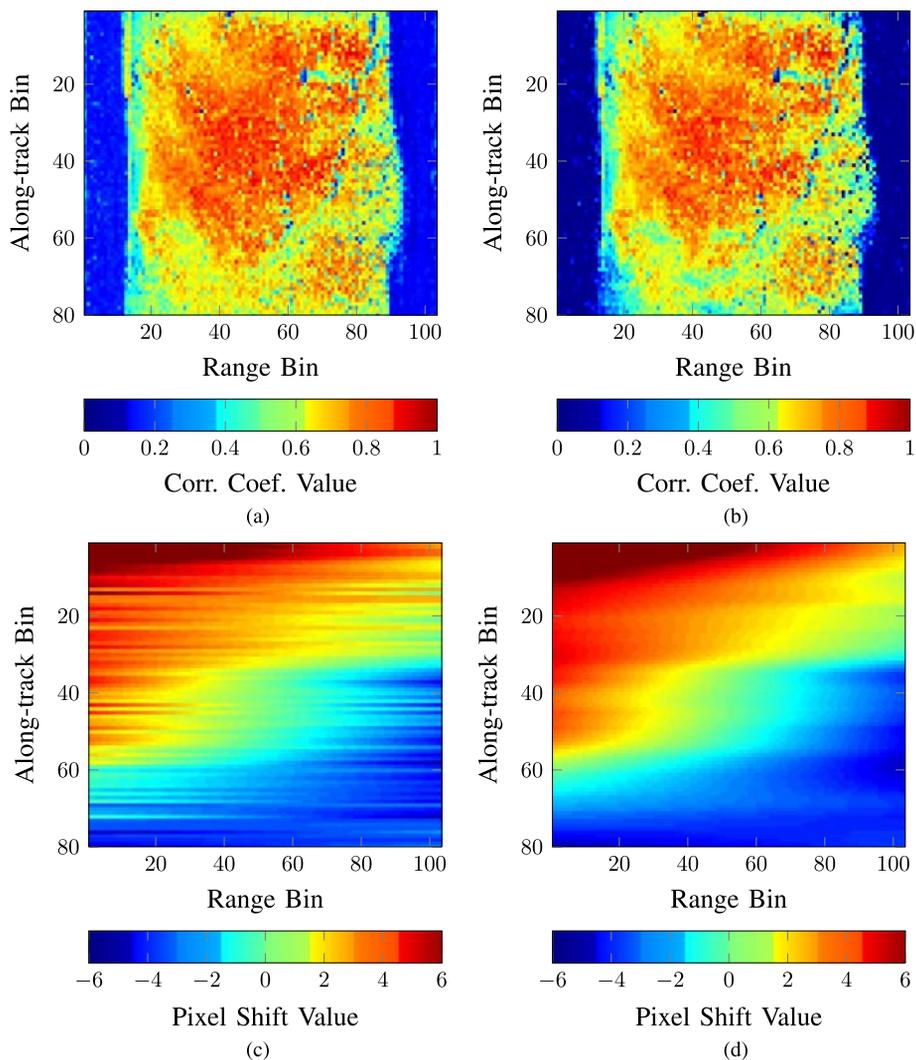


Fig. 8. Model weights and model fit after iterative refinement. (a) Correlation coefficients used as the initial weights for model fitting. After three iterations the weights have been refined to those shown in (b). (c) and (d) Enhancement of model fit quality to the translational shift estimates when the original weighting scheme is used versus the refined weighting scheme.

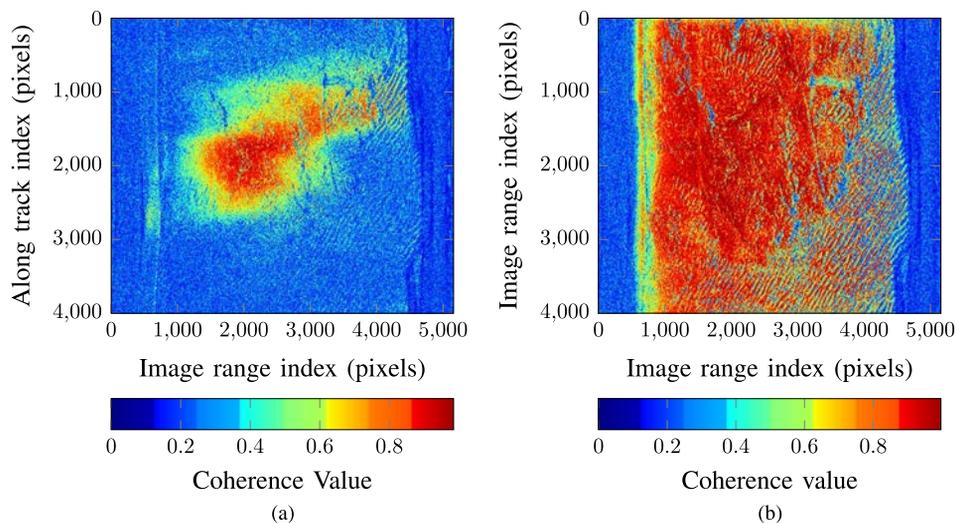


Fig. 9. Local coherence estimates for the reference and matching images (a) before and (b) after the remapping operation to the reference and repeat-pass images shown in Fig. 3. The coherence is computed by applying (17) to a sliding 10×10 pixel neighborhood and calculating the magnitude of the result.

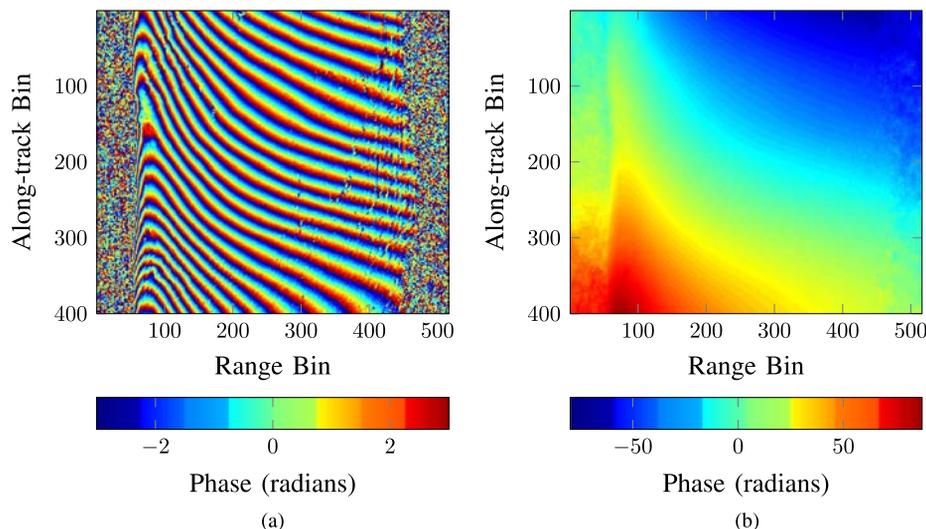


Fig. 10. Images showing (a) interferogram (wrapped phase) and (b) unwrapped phase. Interferograms of the phase are shown in (a), where the color map wraps around every 2π radians.

where the CCD is relying on its ability to effectively remove very subtle changes in the sea surface environment.

1) *Robustness of Change Detection Algorithm:* 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 [see (20)] over data sets collected at site 2. The duration of the data collection was approximately 72 h over this predominantly flat sand environment. To illustrate the coherence measure, Fig. 18(a) and (b) presents box plots of the coherence distribution of SAS images generated before and after phase correction. The coherence was computed between reference and repeat-pass data sets that are coregistered using the automated technique. It should be noted that the values of the CCA in Fig. 18 are the coherence of the entire image; no blocking of the image is done. The abscissae in Fig. 18 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 75th and

25th percentiles. The whiskers are extended to 1.5 times the difference between these percentiles (the interquartile 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 the whiskers represent individual outliers. The total number of samples for the seven repeat-pass data sets used for the formation of the statistics was 247, 251, 249, 320, 316, 246, and 22, respectively. The coherence box plot obtained after the application of local coregistration (phase correction) to the data set, displayed in Fig. 18(b), shows a vast improvement of coherence. We see that the first three repeat-pass data sets significantly improved the coherence except for a few 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, which is caused by currents, marine life, sand ripple formation, and deformation. The first five repeat passes were col-

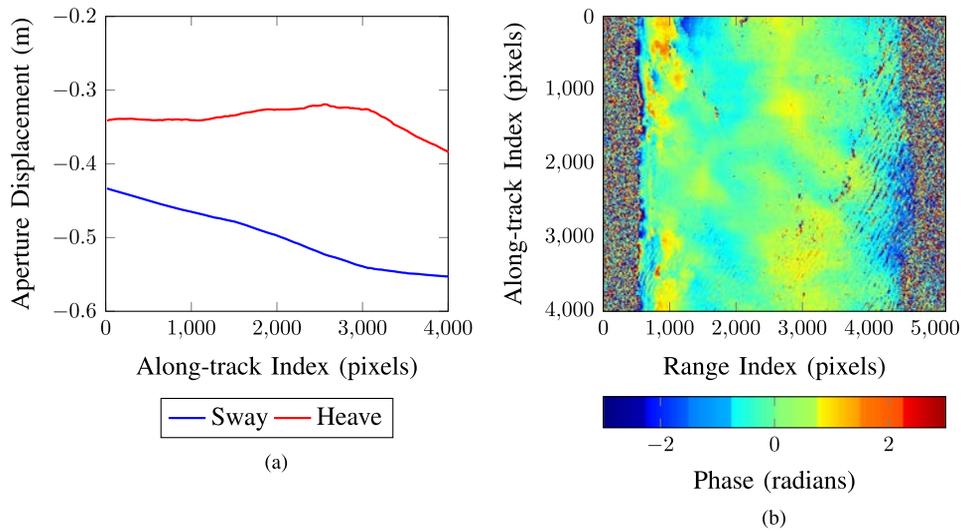


Fig. 11. (a) Estimated heave and sway contributions to the repeat-pass interferometric phase. Heave (vertical) displacements and sway (ground-plane range) displacements are depicted in meters. (b) Interferogram with the heave-and-sway-induced phase-trends removed. Bathymetric features of the sediment are strongly visible at near ranges.

TABLE I

ROBUSTNESS TESTING OF COREGISTRATION FOR DATA SETS. IMAGES FOR THE REFERENCE PASS WERE SHIFTED BY 20, 50, 100, 150, AND 200 PIXELS DOWN AND TO THE RIGHT AND ZERO MEAN GAUSSIAN NOISE WAS ADDED TO FORM REPEAT-PASS DATA FOR COREGISTRATION. THE STANDARD DEVIATION OF NOISE CONSIDERED WAS 1, 1.5, AND 2 TIMES THE IMAGE VARIANCE σ . TABLE VALUES INDICATE THE EUCLIDEAN DISTANCE FROM THE SHIFT ESTIMATED BY THE COREGISTRATION ALGORITHM TO THE TRUE SHIFT

	Noise	Image Shifts (Pixel)				
		20	50	100	150	200
data set 1	σ	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	$1.5 \cdot \sigma$	0.04 ± 0.22	0.07 ± 0.25	0.05 ± 0.22	0.05 ± 0.22	0.04 ± 0.22
	$2.0 \cdot \sigma$	81.57 ± 531.74	84.55 ± 531.68	96.00 ± 536.87	90.27 ± 532.49	85.60 ± 535.34
data set 2	σ	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	$1.5 \cdot \sigma$	0.01 ± 0.09	0.02 ± 0.13	0.01 ± 0.11	0.01 ± 0.09	0.01 ± 0.09
	$2.0 \cdot \sigma$	87.26 ± 252.81	76.32 ± 201.14	81.00 ± 195.92	59.71 ± 159.83	67.36 ± 285.99

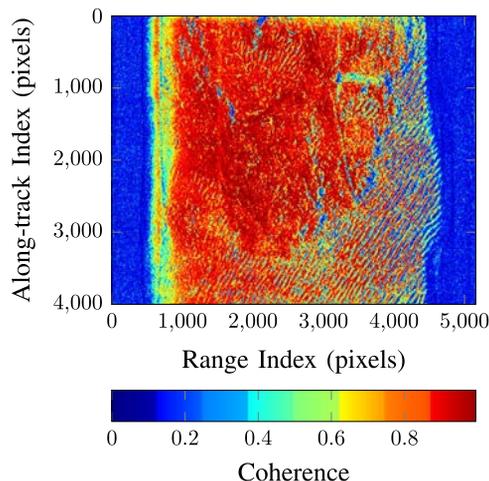


Fig. 12. Coherence of the remapped, phase detrended image computed using a sliding 15×15 pixel window. Comparing with Fig. 9(b), this figure shows that the phase detrending has only a minor effect on the coherence in the present scenario (most values are already near 1, the maximum).

lected along 1-h intervals. The sixth and seventh repeat passes were collected approximately 72 h later.

We determined the percentage of strong correlations and defined strong coherence as a value greater than 60%. The percentage of strong correlations should be regarded as an indication of the ability to perform image-based coherent change detection. The box plot in Fig. 18(a) shows that the data are 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 before running the data through the change detection algorithm should mitigate this problem. Nevertheless, application of local coregistration is a good mitigating factor as shown in Fig. 18(b). The general trend shown in Fig. 18, where the coherence decays (decorrelates) over time, generally agrees with the observations of Lyons and Brown [27], 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 data sets in Fig. 18(a) and (b) represent data collection approximately 72 h after the reference pass. They exhibit a much lower coherence, indicating that this sandy environment,

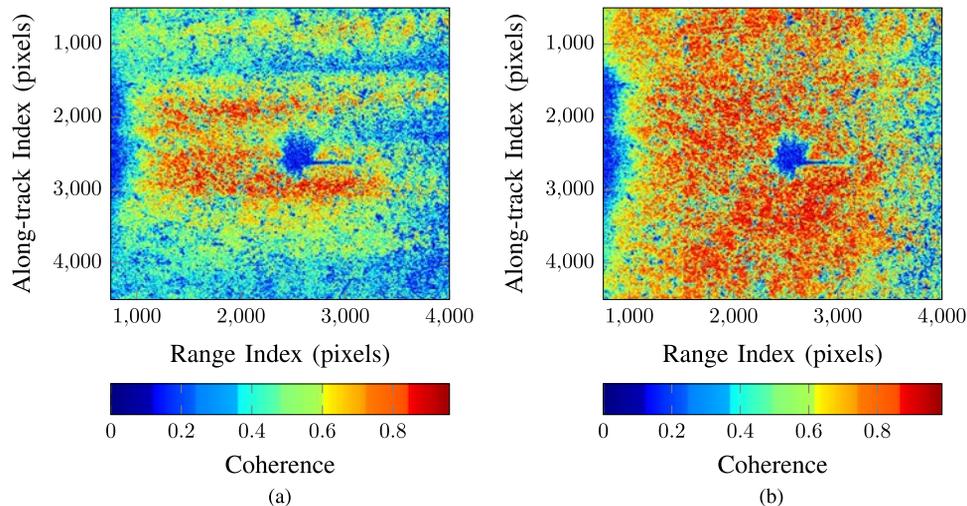


Fig. 13. Correlation coefficient map resulting from application of the local coregistration algorithm (a) without and (b) with application of autofocusing to the complex imagery before coregistration.

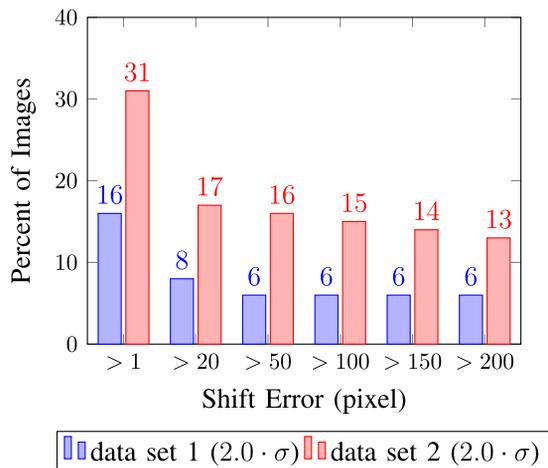


Fig. 14. Histogram of the percentage of images with estimated shift greater than the true shift values for the high noise ($2.0 \cdot \sigma$) scenarios considered in Table I. Histogram is cumulative. For example, all the images with estimated shift > 150 also have an estimated shift > 100 .

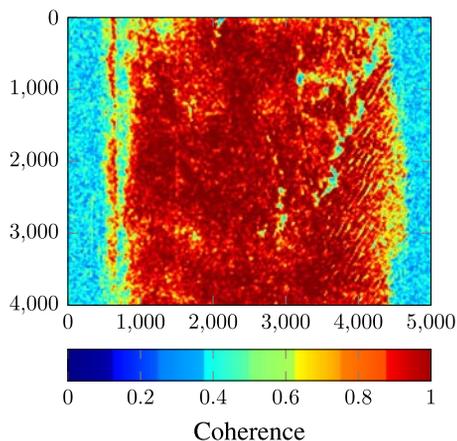


Fig. 15. Coherence maps for coregistered images shown in Fig. 3 using CCA. As in Fig. 12, red indicates high coherence and blue indicates low coherence.

lacking stable features, presents a challenge for change detection over large time scales.

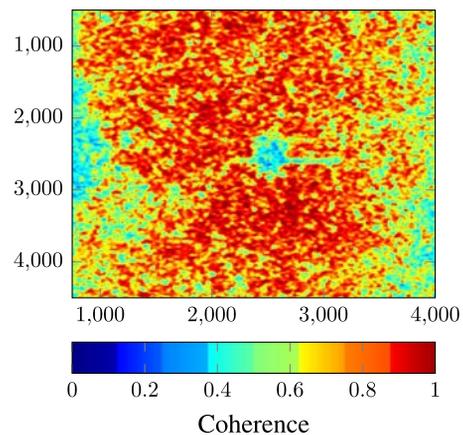


Fig. 16. Coherence maps for coregistered images shown in Fig. 4 using CCA. The coherence map range in value from 0 to 1, with 1 or red indicating high coherence and 0 or blue indicating low coherence.

VII. CONCLUSION AND DISCUSSION

In this paper, we presented and demonstrated a new, robust, and precise seabed change detection method for SAS that automatically compares seafloor images and detects changes over time. Our study primarily focused on the automation of the scene coregistration and coherent change detection. The ACD method is a four-stage process: 1) coarse navigational alignment that relates and approximates the pixel location of reference and repeat-pass data sets; 2) fine-scale coregistration using the SIFT algorithm to match features between overlapping data sets; 3) local coregistration that corrects local discrepancies due to platform motion; and 4) coherent change detection utilizing the CCA algorithm to detect changes. The ACD method was tested using data collected with a high-frequency SAS in a sandy shallow-water environment, subject to currents and tidal changes. By using precise coregistration 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.

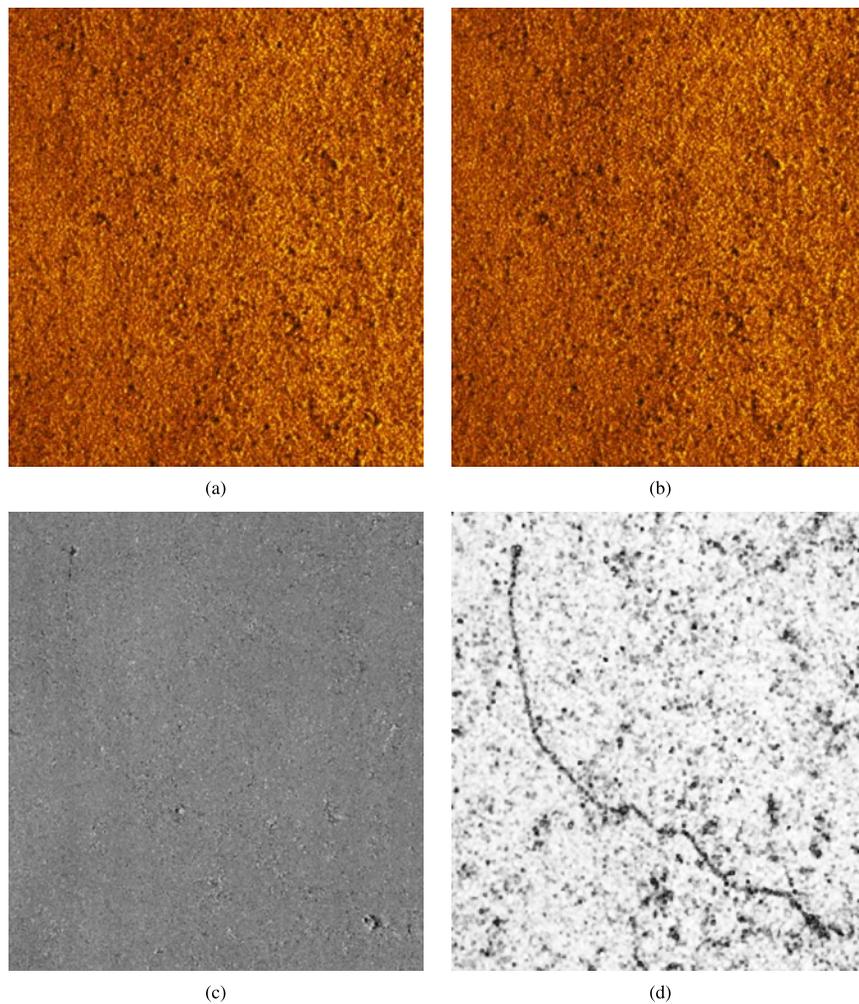


Fig. 17. Change detection comparison between incoherent and coherent change detection. The first pair of images (a) and (b) are snippets of the reference and repeat-pass survey. (c) The incoherent change detection (the despeckled log ratio) image. (d) Coherent change detection (the complex coherence map), where the phase is being used. The CCD clearly shows a strong signal of some sort of object that is not detected in the ICD.

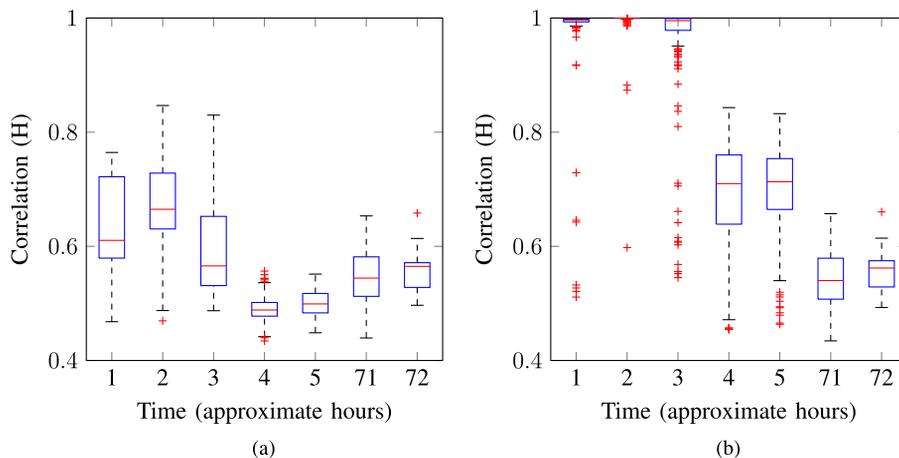


Fig. 18. Coherence box plot. Sample box plots of the measured coherence of seafloor acoustic backscatter H over time. Seven repeat-pass data sets are compared to the reference data set with latencies ranging from 1 to 72 h. (a) Coherence over time before local coregistration. (b) Coherence over time after local coregistration. (a) Before phase correction. (b) After phase correction.

Robustness of the fine-scale coregistration 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 coregister 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 data sets collected with different time latencies showed that coherence between image pairs that were coregistered with the fine-scale SIFT algorithm degraded over time. However, the

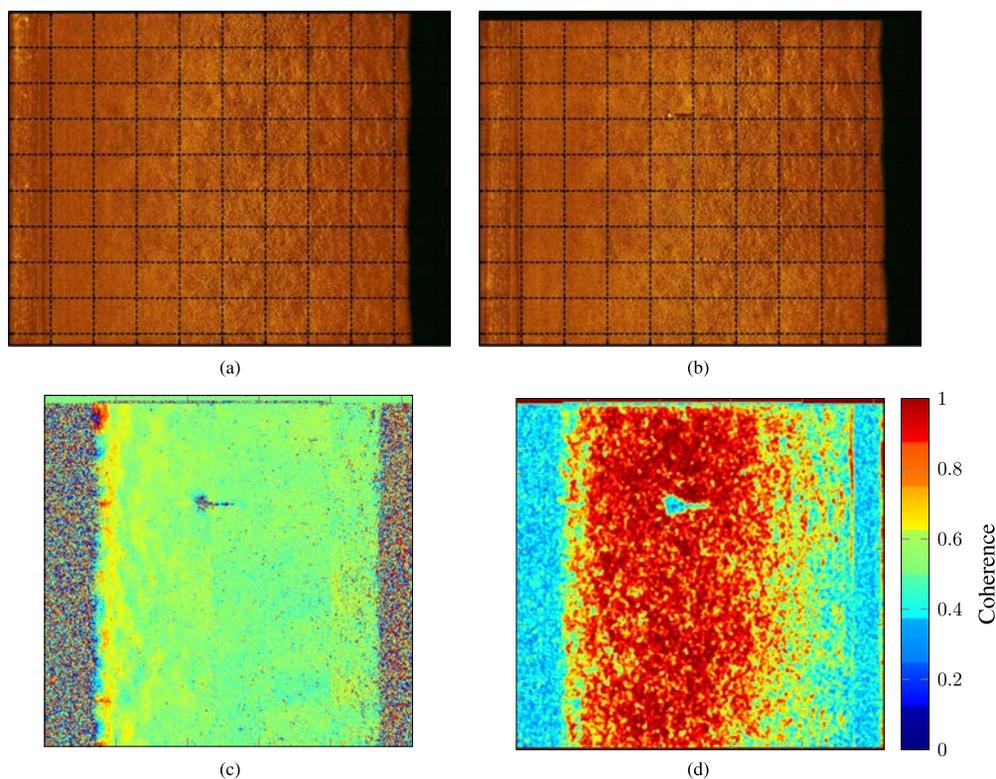


Fig. 19. Example of automated coregistration results: (a) reference and (b) repeat pass. (c) Interferogram and (d) coherence map using CCA are shown for coregistered images.

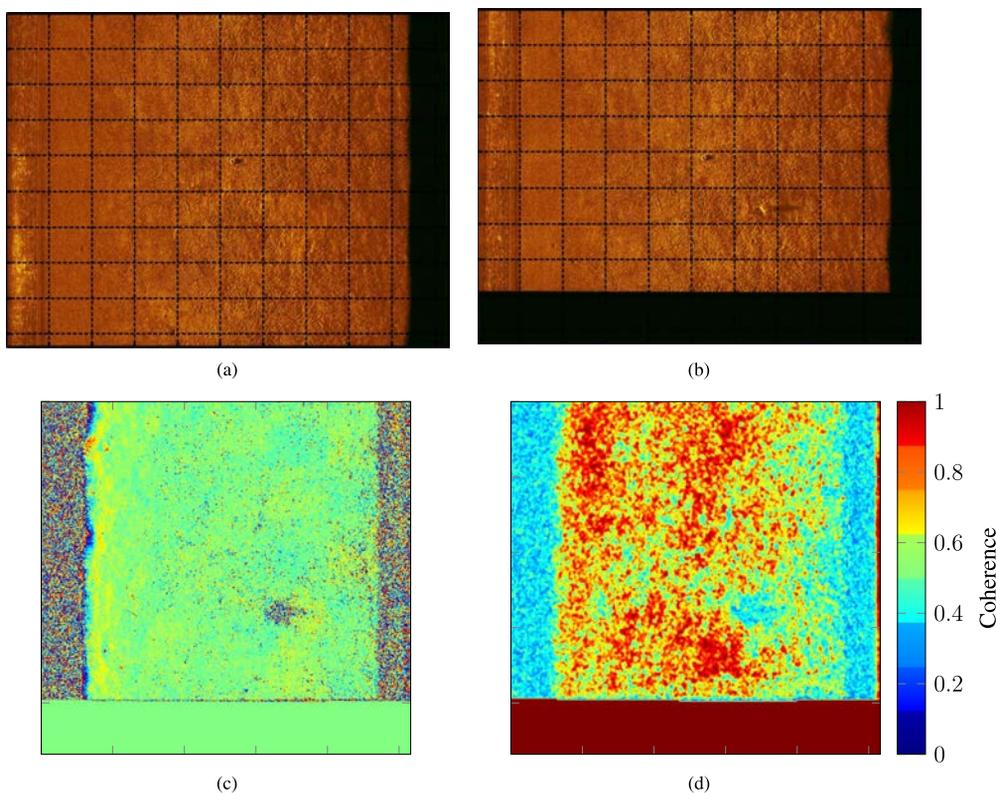


Fig. 20. Example of automated coregistration results: (a) reference and (b) repeat pass. (c) Interferogram and (d) coherence map using CCA are shown for coregistered images.

local coregistration algorithm improves the coregistration, resulting in better interscene coherence.

While the ability to detect changes in signal phase degrades rapidly over time, with time scales dependent on the environ-

ment, the capability of using image-based tools such as those to conduct incoherent change detection for time scales on the order of months has been demonstrated [19]. One benefit of utilizing the phase component is the possibility of detecting a physical

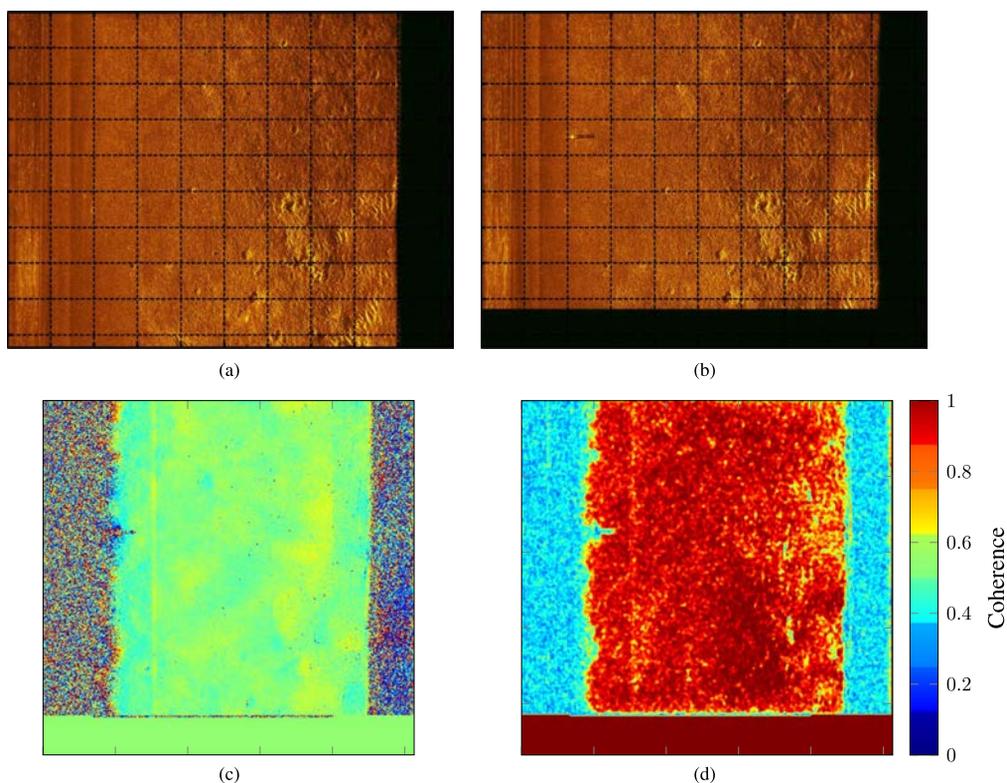


Fig. 21. Example of automated coregistration results: (a) reference and (b) repeat pass. (c) Interferogram and (d) coherence map using CCA are shown for coregistered images.

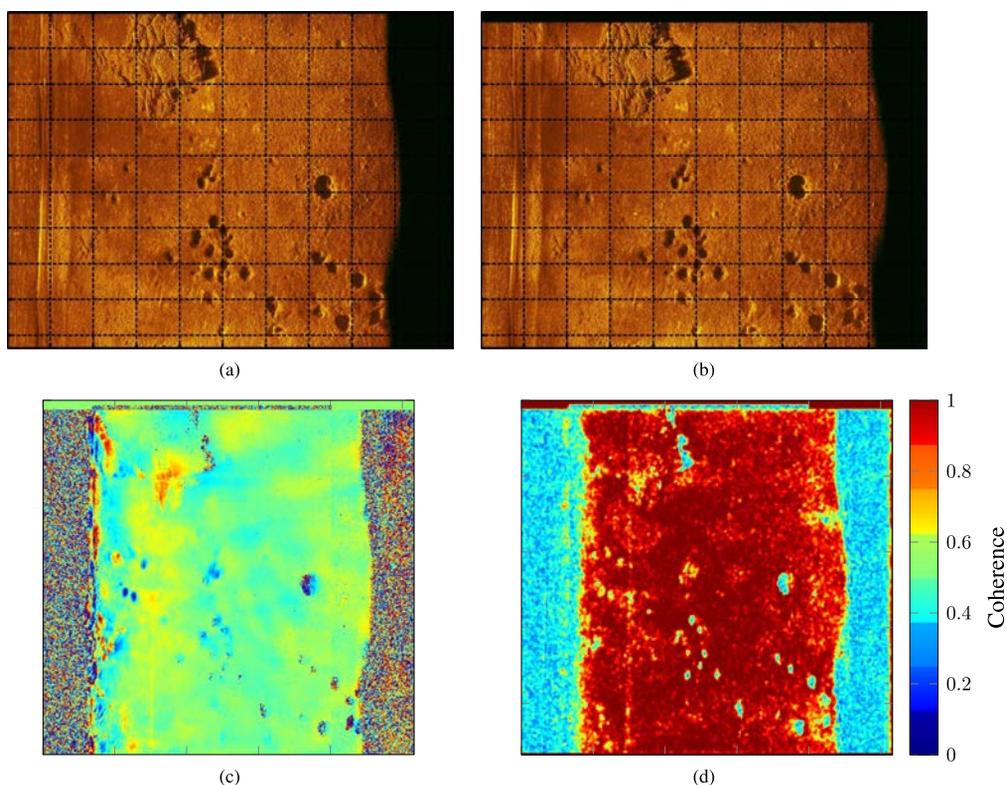


Fig. 22. Example of automated coregistration results: (a) reference and (b) repeat pass. (c) Interferogram and (d) coherence map using CCA are shown for coregistered images.

disturbance in the scene even when there is negligible change in the amplitude of backscattered energy, such as due to a slight rise or fall in topography, due to tectonic disturbance, or due to the introduction of an object onto the seabed that has similar

scattering properties to that of the background. We have demonstrated that coherent ACD is possible for modest time scales and is capable of detecting small disturbances which otherwise would be missed by incoherent methods. Figs. 19–24 are exam-

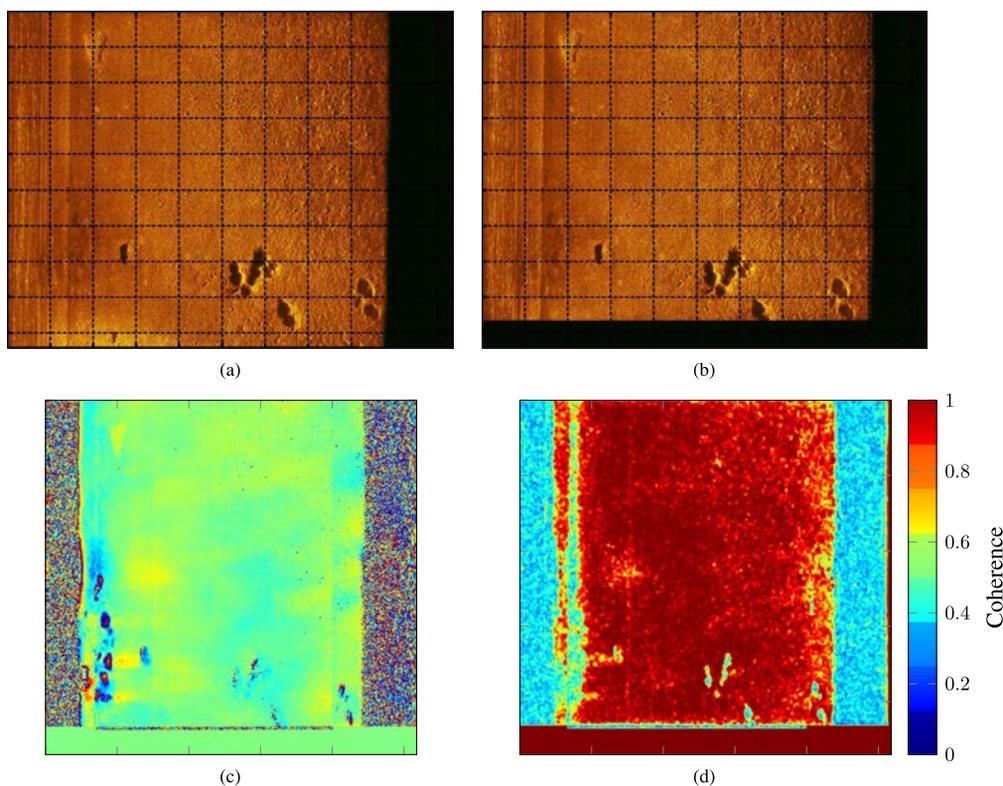


Fig. 23. Example of automated coregistration results: (a) reference and (b) repeat pass. (c) Interferogram and (d) coherence map using CCA are shown for coregistered images.

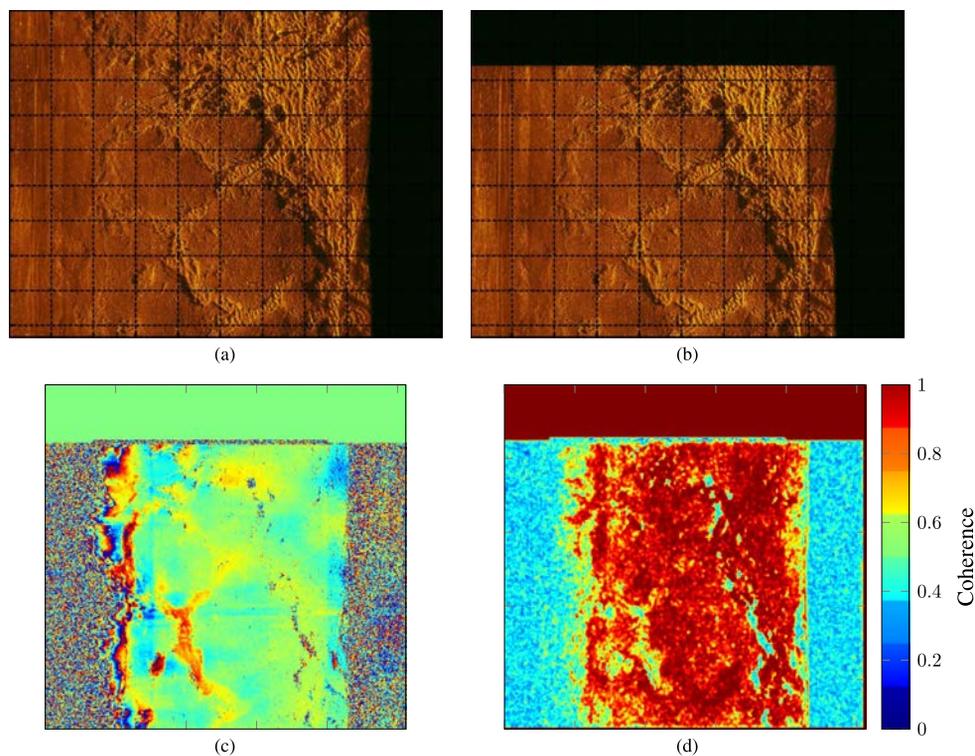


Fig. 24. Example of automated coregistration results: (a) reference and (b) repeat pass. (c) Interferogram and (d) coherence map using CCA are shown for coregistered images.

ples illustrating results of using the ACD method for different background of seafloor.

Further investigation should be carried out to explore the efficacy of utilizing the coherent ACD techniques for detecting small disturbances, over a variety of environments and acoustic

wavelengths. Additionally, a comprehensive comparison of the CCA approach with other coherent change detection methods should be carried out. Other research directions 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 leading toward real-time implementation.

APPENDIX REVIEW OF CCA

CCA was developed by Hotelling as a procedure for assessing the relationship between two data sets [43]. As the name implies, “canonical” refers to mapping the data to a coordinate system that is ideal for capturing linear dependence between the two data channels. That is, this coordinate system reveals the relationships between the two data sets of variables optimally from a coherence (or linear dependence) point of view. The language and terminology used in this section are taken mostly from [40] and [41].

Consider the composite data vector \mathbf{z} consisting of two random vectors $\mathbf{x} \in \mathbb{C}^m$ and $\mathbf{y} \in \mathbb{C}^{mn}$, i.e.,

$$\mathbf{z} = \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \in \mathbb{C}^{(m+n)}. \quad (22)$$

For the remainder of the derivations, it is assumed that $m \geq n$; also the notation $(\cdot)^H$ represents the Hermitian operation. Assume that \mathbf{x} and \mathbf{y} have zero means and share the composite covariance matrix

$$\begin{aligned} R_{zz} &= E[\mathbf{z}\mathbf{z}^H] = E \left[\begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} \begin{pmatrix} \mathbf{x}^H & \mathbf{y}^H \end{pmatrix} \right] \\ &= \begin{bmatrix} R_{xx} & R_{xy} \\ R_{yx} & R_{yy} \end{bmatrix}. \end{aligned} \quad (23)$$

If \mathbf{x} and \mathbf{y} are now replaced by their corresponding whitened vectors, then the composite vector $\boldsymbol{\xi}$ is

$$\boldsymbol{\xi} = \begin{bmatrix} \boldsymbol{\zeta} \\ \boldsymbol{\nu} \end{bmatrix} = \begin{bmatrix} R_{xx}^{-1/2} & \mathbf{0} \\ \mathbf{0} & R_{yy}^{-1/2} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \quad (24)$$

where $R_{xx}^{1/2}$ is a square root of R_{xx} with $R_{xx}^{1/2} R_{xx}^{H/2} = R_{xx}$ and $R_{xx}^{-1/2} R_{xx}^{-H/2} = I$. The covariance matrix of $\boldsymbol{\xi}$ may be written as

$$\begin{aligned} R_{\xi\xi} &= E[\boldsymbol{\xi}\boldsymbol{\xi}^H] = E \left[\begin{pmatrix} \boldsymbol{\zeta} \\ \boldsymbol{\nu} \end{pmatrix} \begin{pmatrix} \boldsymbol{\zeta}^H & \boldsymbol{\nu}^H \end{pmatrix} \right] \\ &= \begin{bmatrix} R_{\zeta\zeta} & R_{\zeta\nu} \\ R_{\nu\zeta} & R_{\nu\nu} \end{bmatrix} \\ &= \begin{bmatrix} I & C \\ C^H & I \end{bmatrix} \end{aligned} \quad (25)$$

where

$$\begin{aligned} C &= E[\boldsymbol{\zeta}\boldsymbol{\nu}^H] = E \left[\left(R_{xx}^{-1/2} \mathbf{x} \right) \left(R_{yy}^{-1/2} \mathbf{y} \right)^H \right] \\ &= R_{xx}^{-1/2} R_{xy} R_{yy}^{-H/2} \end{aligned} \quad (26)$$

is called the coherence matrix of \mathbf{x} and \mathbf{y} [40]. Therefore, the coherence matrix C is the cross-covariance matrix between the whitened versions of \mathbf{x} and \mathbf{y} . Correspondingly, the coordinates $\boldsymbol{\zeta}$ and $\boldsymbol{\nu}$ are called the coherence coordinates. Now it is possible

to determine the singular value decomposition (SVD) of the coherence matrix, namely

$$\begin{aligned} C &= R_{xx}^{-1/2} R_{xy} R_{yy}^{-H/2} = F K G^H \\ F^H C G &= F^H R_{xx}^{-1/2} R_{xy} R_{yy}^{-H/2} G = K \end{aligned} \quad (27)$$

where $F \in \mathbb{C}^{m \times m}$ and $G \in \mathbb{C}^{n \times n}$ are orthogonal matrices [44, Ch. 2.5], i.e.,

$$F^H F = F F^H = I(m) \quad \text{and} \quad G^H G = G G^H = I(n) \quad (28)$$

and

$$K = \begin{bmatrix} K(n) \\ \mathbf{0} \end{bmatrix} \in \mathbb{C}^{m \times n} \quad (29)$$

is a diagonal singular value matrix, with $K(n) = \text{diag}[k_1, k_2, \dots, k_n]$ and $1 \geq k_1 \geq k_2 \geq \dots \geq k_n > 0$.

We then use the orthogonal matrices F and G to transform the whitened composite vector $\boldsymbol{\xi}$ into the canonical composite vector \mathbf{w}

$$\begin{aligned} \mathbf{w} &= \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} = \begin{bmatrix} F^H & \mathbf{0} \\ \mathbf{0} & G^H \end{bmatrix} \begin{bmatrix} \boldsymbol{\zeta} \\ \boldsymbol{\nu} \end{bmatrix} \\ &= \begin{bmatrix} F^H & \mathbf{0} \\ \mathbf{0} & G^H \end{bmatrix} \begin{bmatrix} R_{xx}^{-1/2} & \mathbf{0} \\ \mathbf{0} & R_{yy}^{-1/2} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}. \end{aligned} \quad (30)$$

Then, the covariance matrix for the canonical composite vector \mathbf{w} is obtained as

$$\begin{aligned} R_{ww} &= E[\mathbf{w}\mathbf{w}^H] = E \left[\begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} \begin{pmatrix} \mathbf{u}^H & \mathbf{v}^H \end{pmatrix} \right] \\ &= \begin{bmatrix} R_{uu} & R_{uv} \\ R_{vu} & R_{vv} \end{bmatrix} = \begin{bmatrix} I & K \\ K^H & I \end{bmatrix}. \end{aligned} \quad (31)$$

The elements of $\mathbf{u} = [u_i]_{i=1}^m \in \mathbb{C}^m$ are referred to as the canonical coordinates of \mathbf{x} and the elements of $\mathbf{v} = [v_i]_{i=1}^n \in \mathbb{C}^n$ are the canonical coordinates of \mathbf{y} . The diagonal cross-correlation matrix K

$$\begin{aligned} K &= E[\mathbf{u}\mathbf{v}^H] = E \left[\left(F^H R_{xx}^{-1/2} \mathbf{x} \right) \left(G^H R_{yy}^{-1/2} \mathbf{y} \right)^H \right] \\ &= F^H C G \end{aligned} \quad (32)$$

is called the canonical correlation matrix of canonical correlations k_i , with $1 \geq k_1 \geq k_2 \geq \dots \geq k_n > 0$. Thus, the canonical correlations measure the correlations between pairs of corresponding canonical coordinates. That is, $E[u_i v_j] = k_i \delta_{ij}$; $i \in [1, n]$, $j \in [1, m]$, with δ_{ij} being the Kronecker delta. The canonical correlations k_i are also the singular values of the coherence matrix C . Correspondingly, $K K^H$ is the squared canonical correlation matrix of the squared canonical correlations k_i^2 . Since F and G are orthogonal matrices, we may write the squared coherence matrix $C C^H$ as

$$\begin{aligned} C C^H &= R_{xx}^{-1/2} R_{xy} R_{yy}^{-1} R_{yx} R_{xx}^{-H/2} \\ &= F K G^H G K^H F^H = F K K^H F^H. \end{aligned} \quad (33)$$

This shows that the squared canonical correlations k_i^2 are the eigenvalues of the squared coherence matrix $C C^H$, or equivalently, of the matrix $R_{xx}^{-H/2} C C^H R_{xx}^{H/2} = R_{xx}^{-1} R_{xy} R_{yy}^{-1} R_{yx}$.

It can be shown [40] that the linear dependence measure between the two channels \mathbf{x} and \mathbf{y} is given by

$$L = \det(I - K K^H) = \prod_{i=1}^n (1 - k_i^2), \quad 0 \leq L \leq 1 \quad (34)$$

i.e., the linear dependence is represented in terms of the canonical correlations k_i^2 which measure the dependence between the corresponding canonical coordinates. The measure L takes the value 0 when there is perfect linear dependence between \mathbf{x} and \mathbf{y} , while it takes the value 1 if \mathbf{x} and \mathbf{y} are independent. The i th term of the product on the right-hand side of (34), i.e., $(1 - k_i^2)$, measures the linear dependence between the i th canonical coordinate of \mathbf{x} and the i th canonical coordinate of \mathbf{y} .

The coherence measure between \mathbf{x} and \mathbf{y} is given by

$$H = 1 - L = 1 - \det(I - KK^H) \\ = 1 - \prod_{i=1}^n (1 - k_i^2), \quad 0 \leq H \leq 1. \quad (35)$$

The channels \mathbf{x} and \mathbf{y} are perfectly coherent if $H = 1$, and noncoherent if $H = 0$.

ACKNOWLEDGMENT

The authors would like to thank Dr. J. Stack, Office of Naval Research Code 321OE, Ocean Engineering and Maritime Systems, for guidance and support.

REFERENCES

- [1] J. R. Anderson, "Land use and land cover changes, a framework for monitoring," *J. Res. U.S. Geol. Surv.*, vol. 5, pp. 143–153, Mar.–Apr. 1977.
- [2] A. Singh, "Digital change detection techniques using remotely-sensed data," *Int. J. Remote Sens.*, vol. 10, no. 6, pp. 989–1003, 1989.
- [3] E. J. M. Rignot and J. J. van Zyl, "Change detection techniques for ERS-1 SAR data," *IEEE Trans. Geosci. Remote Sens.*, vol. 31, no. 4, pp. 896–906, Jul. 1993.
- [4] A. Venot and J. L. Golmard *et al.*, "Digital methods for change detection in medical images," in *Proc. 8th Conf. Inf. Process. Medical Imaging*, 1984, pp. 1–16.
- [5] L. Lemieux, U. Wiesmann, 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.*, vol. 2, no. 3, pp. 227–242, Sep. 1998.
- [6] M. Bosc *et al.*, "Automatic change detection in multimodal serial MRI: Application to multiple sclerosis lesion evolution," *NeuroImage*, vol. 20, no. 2, pp. 643–656, Oct. 2003.
- [7] D. Brie, M. Tomczak, H. Oehlmann, and A. Richard, "Gear crack detection by adaptive amplitude and phase demodulation," *Mech. Syst. Signal Process.*, vol. 11, no. 1, pp. 149–167, Jan. 1997.
- [8] R. Bansal and D. Kazakos, "An algorithm for detecting a change in a stochastic process," *IEEE Trans. Inf. Theory*, vol. IT-32, no. 2, pp. 227–235, Mar. 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 Netw.*, vol. 14, no. 3, 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. Pattern Anal. Mach. Intell.*, vol. 32, no. 7, pp. 1239–1258, Jul. 2010.
- [11] E. Coiras, J. Groen, D. Williams, B. Evans, and M. Pinto, "Automatic change detection for the monitoring of cluttered underwater areas," in *Proc. 1st Int. Conf. Exhibit. Waterside Security*, Sep. 2008, pp. 1–6.
- [12] V. Myers, A. Fortin, and P. Simard, "An automated method for change detection in areas of high clutter density using sonar imagery," in *Proc. Underwater Acoust. Meas., Technol. Results*, 2009, pp. 1–8.
- [13] V. Myers, I. Quidu, T. O. Sæbø, and R. E. Hansen, "Results and analysis of coherent change detection experiments using repeat-pass synthetic aperture sonar images," in *Proc. Underwater Acoust. Conf., Cofu, Greece*, Jun. 2013, pp. 519–598.
- [14] P. Gamba, F. Dell'Acqua, and G. Lisini, "Change detection of multitemporal SAR data in urban areas combining feature-based techniques," *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 10, pp. 2820–2827, Oct. 2006.
- [15] M. Gendron, M. Lohrenz, G. Layne, and J. Ioup, "Automatic change detection and classification (ACDC) system," in *Proc. 6th Int. Symp. Technol. Mine Problem*, May 2004, pp. 1–9.
- [16] D. D. Sternlicht, J. K. Harbaugh, and M. A. Nelson, "Experiments in coherent change detection for synthetic aperture sonar," in *Proc. MTS/IEEE OCEANS Conf.*, 2009, pp. 1–5.
- [17] T. G-Michael and J. D. Tucker, "Canonical correlation analysis for coherent change detection in synthetic aperture sonar imagery," *Inst. Acoust. Proc.*, vol. 32, no. 4, pp. 117–122, Sep. 2010.
- [18] C. A. Matthews and D. Sternlicht, "Seabed change detection in challenging environments," *Proc. SPIE—Int. Soc. Opt. Eng.*, vol. 8017, Detection and Sensing of Mines, Explosive Objects, Obscured Targets XVI, May 2011, 80170P-80.
- [19] Ø Midtgaard *et al.*, "Change detection using synthetic aperture sonar: Preliminary results from the larvik trial," in *Proc. MTS/IEEE OCEANS Conf.*, 2011, pp. 1–8.
- [20] I. Quidu, V. Myers, Ø Midtgaard, and R. E. Hansen, "Subpixel image registration for coherent change detection between two high resolution passes," in *Proc. Int. Conf. Underwater Remote Sens.*, Brest, France, Oct. 2012, pp. 207–226.
- [21] Ø. Midtgaard, "Change detection using synthetic aperture sonar imagery with variable time intervals," in *Proc. Underwater Acoust. Conf., Cofu, Greece*, Jun. 2013, pp. 713–720.
- [22] A. V. Oppenheim and J. S. Lim, "The importance of phase in signals," *Proc. IEEE*, vol. 69, no. 5, pp. 529–541, May 1981.
- [23] T. Huang, J. Burnett, and A. Deczky, "The importance of phase in image processing filters," *IEEE Trans. Acoust. Speech Signal Process.*, vol. 23, no. 6, pp. 529–542, Jun. 1975.
- [24] M. Preiss and N. J. S. Stacey, "Coherent change detection: Theoretical description and experimental results," Tech. Rep. DSTO-TR-1851, Aug. 2006.
- [25] D. R. Jackson *et al.*, "Acoustic observation of the time dependence of the roughness of sandy seafloors," *IEEE J. Ocean. Eng.*, vol. 34, no. 4, pp. 407–422, Oct. 2009.
- [26] A. L. Gerig, A. P. Lyons, E. Pouliquen, and K. L. Williams, "Comparison of seafloor roughness and scattered acoustic temporal decorrelation," *IEEE J. Ocean. Eng.*, vol. 34, no. 4, pp. 423–430, Oct. 2009.
- [27] 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. Ocean. Eng.*, vol. 38, no. 1, pp. 91–97, Jan. 2013.
- [28] M. R. Azimi-Sadjadi and S. Srinivasan, "Coherent change detection and classification in synthetic aperture radar imagery using canonical correlation analysis," Internal Rep., Information System Technologies, Inc., Apr. 2007.
- [29] D. D. Sternlicht and T. G-Michael, "Change detection by image correlation for synthetic aperture sonar," *Inst. Acoust. Proc.*, vol. 32, no. 4, pp. 87–91, Sep. 2010.
- [30] J. D. Tucker and M. R. Azimi-Sadjadi, "Coherence-based underwater target detection from multiple sonar platforms," *IEEE J. Ocean. Eng.*, vol. 36, no. 1, pp. 37–51, Jan. 2011.
- [31] D. G. Lowe, "Object recognition from local scale-invariant features," in *Proc. 7th IEEE Int. Conf. Comput. Vis.*, 1999, vol. 2, pp. 1150–1157.
- [32] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *Int. J. Comput. Vis.*, vol. 60, no. 2, pp. 91–110, Nov. 2004.
- [33] D. D. Sternlicht *et al.*, "Advanced sonar technologies for autonomous mine countermeasures," in *Proc. MTS/IEEE OCEANS Conf.*, 2011, pp. 1–5.
- [34] A. Vedaldi and B. Fulkerson, "VLFeat: An open and portable library of computer vision algorithms," 2008 [Online]. Available: <http://www.vlfeat.org/>
- [35] G. Llort-Pujol, C. Sintes, T. Chonavel, I. A. T. Morrison, and S. Daniel, "Advanced interferometric techniques for high-resolution bathymetry," *J. Mar. Technol. Soc.*, vol. 46, no. 2, pp. 9–31, Mar. 2012.
- [36] T. Sæbø, R. E. Hansen, H. J. Callow, and S. A. V. Synnes, "Coregistration of synthetic aperture sonar images from repeated passes," in *Proc. Underwater Acoust. Meas., Technol. Results*, 2011, pp. 1–8.
- [37] W. G. Carrara, R. S. Goodman, and R. M. Majewski, *Spotlight Synthetic Aperture Radar Signal Processing Algorithms*. Norwood, MA, USA: Artech House, 1995.
- [38] H. J. Callow, M. P. Hayes, and P. T. Gough, "Autofocus," in *Proc. OCEANS Conf.*, 2003, vol. 5, pp. 2422–2426.
- [39] D. C. Ghiglia and M. D. Pritt, *Two-Dimensional Phase Unwrapping: Theory, Algorithms and Software*, 1st ed. ed. New York, NY, USA: Wiley-Interscience, Apr. 1998, ch. 5.

- [40] L. Scharf and C. Mullis, "Canonical coordinates and the geometry of inference, rate, capacity," *IEEE Trans. Signal Process.*, vol. 48, no. 3, pp. 824–891, Mar. 2000.
- [41] 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.*, vol. 54, no. 12, pp. 4817–4820, Dec. 2006.
- [42] H. J. Callow, "Comparison of SAS processing strategies for crabbing collection geometries," in *Proc. MTS/IEEE OCEANS Conf.*, 2010, DOI: 10.1109/OCEANS.2010.5664446.
- [43] H. Hotelling, "Relations between two sets of variates," *Biometrika*, vol. 28, pp. 321–377, 1936.
- [44] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. Baltimore, MD, USA: Johns Hopkins Univ. Press, 1996.



Tesfaye G-Michael (M'96–SM'12) received the B.S. degree in biomedical engineering, the M.S. degree in electrical engineering, and the M.S. degree in computer engineering from Wright State University, Dayton, OH, USA, in 1996, 1998, and 2002, respectively.

From 1997 to 2002, he was an Engineer at Jacobs Sverdrup Engineering–Advanced Systems Group, where he was involved in a variety of projects in Automatic Target Detection and evaluation, radar system design, and simulation. From 2002 to 2006,

he was a Senior Radar Engineer at MacAuly Brown Inc., Integrated Avionics Test Facility (IATF). From 2006 to 2008, he was Senior System Engineer II at Raytheon company. Currently, he is a Research Engineer at Naval Surface Warfare Center Panama City Division, Panama City, FL, USA, where he is leading the Automated Change Detection project. His research interest includes synthetic aperture sonar (SAS) and synthetic aperture radar (SAR) target detection, classification and tracking, image and signal processing, and bistatic/multistatic SAR.

Mr. G-Michael is a Senior Member of the American Institute of Aeronautics & Astronautics (AIAA).



Bradley Marchand received the B.S. degree in mathematics from California State University, Bakersfield, CA, USA, in 2004 and the Ph.D. degree in applied mathematics from the University of California, Davis, CA, USA, in 2010.

Since April 2010, he has been with the Naval Surface Warfare Center Panama City Division (NSWC PCD), Panama City, FL, USA. His research interests include harmonic analysis, feature extraction, pattern recognition, classification, and machine learning.



J. Derek Tucker (S'04–M'14) received the B.S. and M.S. degrees in electrical engineering, with an emphasis on statistical signal processing, from Colorado State University, Fort Collins, CO, USA, in 2007 and 2009, respectively, and the Ph.D. degree in statistics from Florida State University, Tallahassee, FL, USA, in 2014.

He is currently with Sandia National Laboratories, Albuquerque, NM, USA, as a senior Member of the Technical Staff in the Radar and Signal Analysis Group. His research interests include statistical

image understanding, target detection, classification, functional data analysis, and statistical learning methods.



Timothy M. Marston received the B.S. degree in electrical engineering from Seattle Pacific University, Seattle, WA, USA, in 2004 and the M.S. and Ph.D. degrees in acoustics from the Pennsylvania State University, State College, PA, USA, in 2006 and 2009, respectively.

In 2010, he joined the Naval Surface Warfare Center Panama City Division (NSWC PCD), Panama City, FL, USA, and his research primary focus was the development of robust algorithms for synthetic aperture data processing. He is currently

employed by the Acoustics Department, Applied Physics Laboratory (APL), University of Washington, Seattle, WA, USA.



Daniel D. Sternlicht (M'88) received the B.A. degree in biology from the University of Pennsylvania, Philadelphia, PA, USA, in 1983, the M.S. degree in electrical engineering from the University of Hawaii at Manoa, Honolulu, HI, USA, in 1992, and the Ph.D. degree in electrical engineering and applied ocean science from the University of California, San Diego, La Jolla, CA, USA, in 1999.

From 1998 to 2003, he was Senior Principal Engineer and Countermeasures Program Manager at Lockheed Martin ORINCON; from 2003 to 2008, he

was Director of Synthetic Aperture Sonar Programs at Dynamics Technology Inc. (now Raytheon Applied Signal Technology). He is currently the head of the Sensing Sciences Division at the Naval Surface Warfare Center Panama City Division, Panama City, FL, USA, which develops advanced sensors and processing for U.S. Navy and Marine Corps missions.

Dr. Sternlicht chaired the San Diego Chapter of IEEE Oceanic Engineering Society (2001–2003), served on the OCEANS 2003 Technical Program Committee, served on the IEEE Oceanic Engineering Society Administrative Committee (2004–2007), and was a guest editor for the 2009 IEEE JOURNAL OF OCEANIC ENGINEERING Special Issue on Synthetic Aperture Sonar. He is a member of the Acoustical Society of America, the National Defense Industrial Association, and the Mine Warfare Association.



Mahmood R. Azimi-Sadjadi (S'81–M'81–SM'89) received the M.S. and Ph.D. degrees in electrical engineering with specialization in digital signal/image processing from the Imperial College of Science and Technology, University of London, London, U.K., in 1978 and 1982, respectively.

Currently, he is a Full Professor at the Electrical and Computer Engineering Department, Colorado State University (CSU), Fort Collins, CO, USA. He is also the Director of the Digital Signal/Image Laboratory, CSU. His main areas of interest include

digital signal and image processing, wireless sensor networks, target detection, classification and tracking, adaptive filtering, system identification, and neural networks. His research efforts in these areas resulted in over 250 journal and referenced conference publications. He is the coauthor of the book *Digital Filtering in One and Two Dimensions* (New York, NY, USA: Plenum, 1989).

Prof. Azimi-Sadjadi served as an Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING and the IEEE TRANSACTIONS ON NEURAL NETWORKS.