AUTOMATED CHANGE DETECTION IN STREAMING SAS IMAGERY

Ø. Midtgaard Norwegian Defence Research Establishment (FFI), Kjeller, Norway

1 INTRODUCTION

Change Detection (CD) is a processing technique which extracts objects or features of interest as differences between two images of the same scene. It is widely used in a variety of domains with diverse sensor types. Mine hunting is an important naval application of CD. Reference sonar imagery of strategic ports, inlets or sea lines of communication is recorded during route surveys, when the seafloor is assumed to be free of mines. After a new survey, mines are recognized as objects that are only present in the current imagery. CD can be the only feasible option when the seafloor clutter density is high, or when the mine's physical characteristics are unknown, as for improvised explosive devices (IED). Other underwater applications include seabed infrastructure inspection, environmental monitoring, Intelligence, Surveillance and Reconnaissance (ISR) and marine science. Due to the high costs of large area seafloor surveys, typical time intervals may be in the order of one year or more.

As manual analysis of large image sets is tedious and prone to error, there is a need for Automated Change Detection (ACD) systems that can compare imagery and reliably report the relevant changes. ACD approaches can be broadly categorized based on the data level used for temporal matching: decision level methods apply a trained or modelled recognition algorithm on each data set and compare the results, while image level methods compare the data sets directly, either for regions or individual pixels. For coherent sensors, image-level ACD can be further categorized as coherent (utilizing both pixel phase and magnitude) or incoherent (magnitude only). Coherent processing is more sensitive than incoherent to minute scene changes, but imposes stricter operational requirements.

ACD for mine hunting was originally focused on decision level methods based on geographical association of contacts (detected mine-like objects) across the data sets [1,2]. This was due to shortfalls of traditional side-scan sonar (SSS) imagery, e.g. poor and range-dependent along-track resolution, as well as challenges imposed by the underwater environment regarding data positioning, platform trajectory control and signal propagation. The introduction of Autonomous Underwater Vehicles (AUV) equipped with high-grade Aided Inertial Navigation Systems (AINS) and Synthetic Aperture Sonar (SAS) has partly remedied these deficiencies and facilitated development of image level sonar ACD [3-7]. These works have focused on smaller, selected seabed scenes, where a pair of corresponding images has been produced through synchronized SAS processing of old and new data. The technical contribution of this paper is to demonstrate ACD for a full survey in a varied ocean environment, using input images created through independent SAS processing of the two data sets. This constitutes an important development step towards an efficient and operational system, which is the aim of FFI's work.

2 SAS CHANGE DETECTION PROCESSING

2.1 Coherent and incoherent processing

Coherent ACD performs complex cross-correlation for small, co-registered image windows and detects changes as locations with reduced coherence (data similarity). The method is sensitive to the distribution of individual sonar scatterers within a resolution cell and can detect even subtle

changes that are invisible in the magnitude SAS image. This also makes it vulnerable to temporal decorrelation of the scene background, as only minute perturbations caused by e.g. currents or marine life are sufficient to reduce the repeat pass coherence. The feasible interval for coherent repeat pass processing of high-frequency SAS data depends on the ocean environment, but has been estimated to maximum a few days [8]. Additionally, the method requires data co-registration accuracy within one-tenth of the pixel size and small sensor trajectory offsets to avoid decorrelation due to sensing geometry differences. Finally, high repeat pass coherence can only be achieved when the single pass coherence is high in both input images. Reduced coherence in regions with multipath pollution and low reverberation (e.g. sonar shadows) can then be mistaken for changes, unless specifically addressed by e.g. masked coherence [9].

Contrary, incoherent ACD can only detect changes in sonar image magnitude, but has significantly more relaxed requirements. Most importantly, survey intervals can be considerably longer. Intervals in the order of years have been demonstrated [10], which comply with operational demands for most applications. The needed co-registration accuracy is around the size of a pixel. The method is also more tolerant to sensor trajectory offsets and image degradation. The ability to produce a change image revealing the magnitude differences between the two co-registered input images is another advantage of incoherent ACD. Analysis of this difference image is usually intuitive, as new objects have similar contrast and shape as in the repeat-pass image, while removed objects have inverted contrast compared to the reference image. This simplifies detection of only specific types of change.

As FFI aims to develop a naval mine hunting ACD system, robustness under realistic operating conditions is essential. This study is thus applies incoherent processing.

2.2 Streaming SAS data

Strip-map SAS processing can operate in either streaming or spot mode [11]. Streaming mode produces a continuous, full-range swath of slightly overlapping images and is the preferred mode to generate imagery of a full SAS survey, commonly applying fast wavenumber domain algorithms to reduce processing time. Spot mode is used to generate superior quality images of selected, smaller scenes, applying more processing intensive algorithms like time-domain backprojection. Previous studies on SAS ACD have typically applied spot generated reference and repeat pass image pairs of the same seafloor regions, based on vehicle navigation data. For an operational ACD system, use of the standard SAS survey products yields benefits such as reduced data storage demands and processing time. However, the resulting unsynchronized image blocking introduces another processing challenge as each repeat pass image partly overlaps several reference images (Figure 1). Streaming images from the HISAS 1030 sonar are projected in slant-range and stored in the XTF file format. Succeeding images with approximately similar orientations are concatenated into a single file after overlap removal. Image-wise ACD processing is necessary as image resolution and motion compensated heading vary slightly between image blocks.

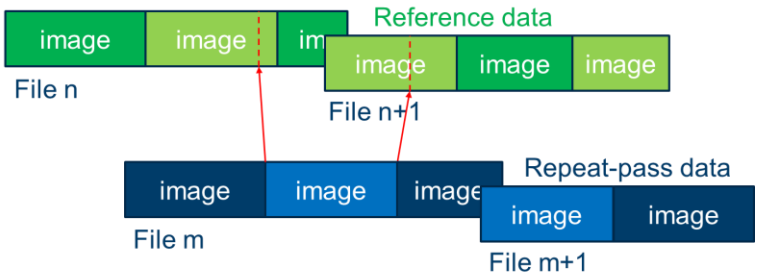

Figure 1: *The HISAS streaming data files contain multiple, adjacent SAS images. There is a few meters image overlap between succeeding files, but no overlap within a file. Each repeat-pass image corresponds to parts of several reference images, which must be sequentially retrieved for ACD processing, as co-registration parameters must be determined separately for each image pair, due to (minute) grid coordinate variations between neighbour images. Margins are added when extracting reference imagery, to account for relative navigation errors.*

2.3 Processing chain

Figure 2 shows the data processing chain used in this study. For each repeat pass image, corresponding reference imagery is sequentially retrieved based on vehicle navigation data (Figure 1) and accurately aligned with the repeat pass image to build a change image mosaic. The pre-processing starts with a logarithmic transform of magnitude to emphasize sonar highlights and shadows more equally. Expected spatial and temporal decorrelation under typical operating conditions will cause image speckle for the two passes to differ. Anisotropic diffusion[12] filtering to reduce speckle while preserving edges, improves the data-driven co-registration and lowers the "noise floor" in the change image. The images are also resampled from slant-range to ground-range coordinates using high-resolution bathymetry from interferometric sonar data processing.

The fine-scale image co-registration is based on feature point matching. Previous SAS ACD studies[4,7] have used the U-SURF[13] or SIFT[14] algorithms to extract and describe feature points in the images. Based on evaluation of state-of-the-art algorithms on varied SAS CD data, however, this work applies the faster CenSurE/STAR [15] detector and BRIEF [16] descriptor. As an example, Table 2 lists feature point results for the two scenes in Figure 5 using four different methods. Point pairs are robustly matched based on description similarity and model fit, and used to estimate the parameters for an affine transformation of the reference image onto the pixel coordinates of the repeat pass image. Pixel-wise subtraction of magnitude values then produces the change image section corresponding to the reference image's overlap with the repeat pass image. All retrieved reference images are processed like this, before proceeding with the next repeat pass image, etc.

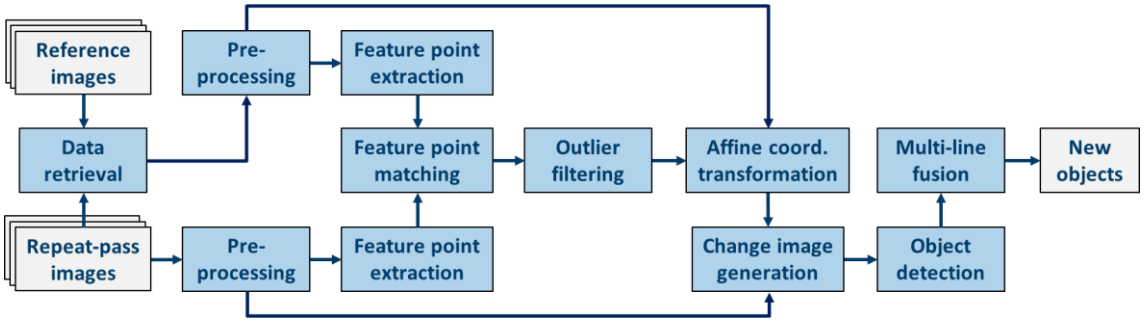

Figure 2: *Processing chain for incoherent, image-level ACD on streaming SAS data.*

For several ACD applications including naval mine hunting, the changes of interest are new objects that are only present in the repeat pass images. However, change images typically reveal a variety of other differences caused by sensing geometry variations, image degradation, data co-registration errors and irrelevant temporal changes due to marine life, currents, waves, human activities, etc. Simple change image thresholding will then not suffice. To detect the objects of interest without a significant number of false alarms, the processing needs to take the specific characteristics of the target changes into account. The matched filter[17] has proven to be a simple, yet effective and versatile detector for mines in sonar images. The filter mask consists of two parts, corresponding to a generic mine highlight region with a trailing shadow region. The length of the shadow region reflects the expected target height and thus increases with range. The mask is convolved with the full change image. In this work, the filter output at each pixel location is the ratio of the average values within the highlight and shadow regions, respectively. The filter is sensitive to proud, mine-sized objects on the seafloor, while suppressing other responses in the image. A classifier assigns a confidence value to each detected change, based on response size and contrast measures.

Seabed locations may be imaged several times during a survey due to sonar coverage overlap. Change detection results can then be fused across survey lines to remove inconsistent detections due to fish, co-registration errors, etc. First, detections corresponding to the same stationary object are clustered, based on geographical positions. Then, any additional sonar views of these locations are added to the clusters as missing detections. Dempster-Shafer Theory of Evidence [18] is used to

combine the positive (detections) and negative (missing detections) evidence into a fused confidence value for each potential object. In this work, objects with fused confidence values exceeding 0.5 and at least two single-view detections were reported as detected changes.

3 SEA TRIALS

The MANEX'14 exercise was organized by the NATO Science and Technology Organisation Centre for Maritime Research and Experimentation (CMRE) and was conducted from the NATO research vessel *Alliance* off the Ligurian coast in Italy in September 2014. FFI participated with scientists and the HUGIN-HUS AUV equipped with HISAS 1030. This interferometric SAS has 100 kHz centre frequency and 30 kHz bandwidth. The typical streaming image resolution is 4cm x 4cm.

Figure 3 shows HISAS bathymetry and sonar mosaic of the operation area in Bonassola bay. The water depths range from 11 to 56 meters. Close to shore the seafloor is varied, with distinct regions of sand ripples, sea grass (*posedonia oceanica*), soft sediments, rocks and rock crops. The deeper water regions contain mostly smooth sediments, but internal waves[19] in the water column introduced significant sonar image texture. A reference HISAS survey was performed on September 23, followed by the deployment of multiple targets on the seafloor. The reference and repeat surveys were intended to run similar mission plans, but due to time constraints and technical difficulties only parts of the plan were executed, and then with some track misalignments, as displayed in Figure 4. The two repeat surveys were performed on September 24 in the time interval from 12:40 to 15:20 and on September 26 from 06:50 to 13:20, respectively.

After the reference survey ten targets were deployed on the seafloor distributed along two recovery ropes. This included six dummy mines (CD1: cylinder, MD1-4: four truncated cones, RD1: wedge shape), two mine-sized rocks (Rock1-2) and two dummy diver bottles (Bottle1-2). In addition, an iron weight (Ballast1-2) was attached to the end of each recovery line. The objects Bottle1-2 and Ballast1-2 were smaller than mine size, but the ballasts produced strong highlights for most of the sonar views. In this paper, the ballasts are considered as deployed objects.

The streaming SAS and ACD processing were performed on board the *Alliance* during the trials.

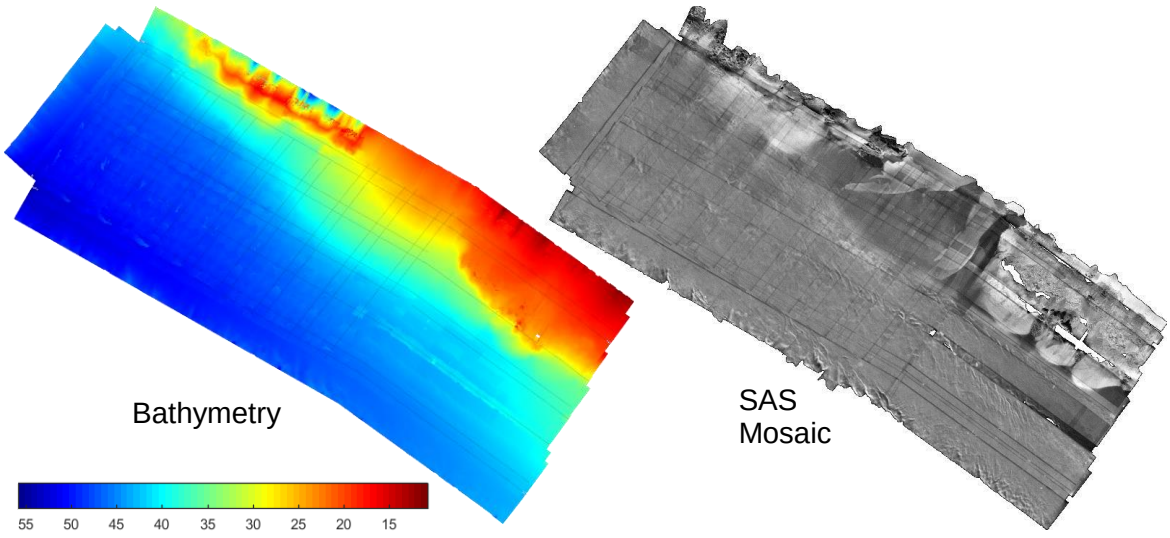

Figure 3: *The survey area in Bonassola bay covered approximately 1.0km x 2.8km, with water depths ranging from 11 to 56 meters (left). The SAS mosaic (right) reveals complex and varied seafloor conditions in the shallow regions, including rock crops, sand ripples and patches of sea grass. The deeper regions have smooth sediment seafloor, but internal waves produced significant sonar image texture.*

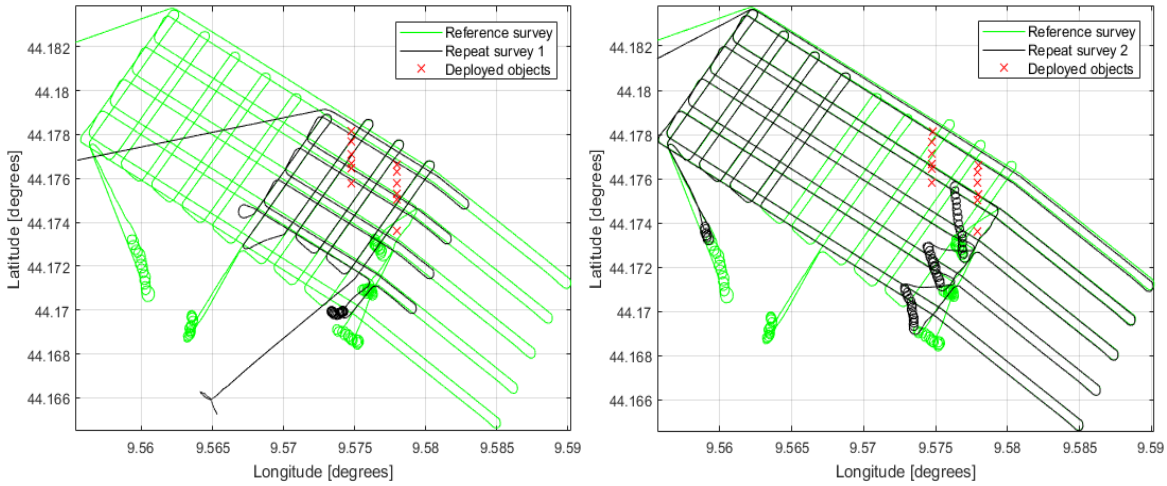

Figure 4: *The vehicle paths (based on post-processed navigation) from the reference survey (green) on September 23 before deployment of seafloor objects (red) and from the two repeat surveys (black) on September 24 (left) and 26 (right). The spirals correspond to vehicle surfacing for GPS position updates.*

4 RESULTS

All streaming SAS images from the surveys were processed through the ACD scheme in Figure 2. The final results were ten detected changes from each repeat survey, as presented in Table 1. The highest fused confidence values correspond to the eight deployed mine-sized targets. The smaller ballast objects were both detected for survey 1, but only Ballast2 was detected for survey 2. The smaller Bottle1-2 objects were not detected. The listed number of object views (detection opportunities) includes all passes with detection of the object and passes without detection if the object lies within the sonar's effective range, as estimated from the interferometric coherence. The table shows that the eight mine-sized objects were detected for all views, except one pass where manual examination reveals that co-registration failed due to marginal overlap between reference and repeat pass images. The total number of object views in Survey 1 is twice that of survey 2 (58 vs. 29), because the objects then was imaged using two perpendicular lawnmower patterns (Figure 4). This may also explain why the fused confidence values seem slightly higher for survey 1.

Repeat survey 1 24-Sep-2014				Repeat survey 2 26-Sep-2014			
Object	**Fused conf.**	**# detect's**	**# views**	**Object**	**Fused conf.**	**# detect's**	**# views**
CD1	1.00	4	4	CD1	1.00	2	2
MD1	1.00	5	6	MD2	1.00	4	4
MD4	1.00	7	7	MD4	1.00	4	4
MD3	1.00	6	6	RD1	1.00	4	4
MD2	1.00	5	5	MD3	0.99	2	2
RD1	1.00	7	7	Rock1	0.99	4	4
Rock2	1.00	6	6	MD1	0.98	2	2
Rock1	1.00	5	5	Rock2	0.96	2	2
Ballast2	0.86	5	6	Ballast2	0.58	2	3
Ballast1	0.66	4	6	Unknown	0.58	2	2

Table 1: *All detected changes from the two repeat surveys after multi-view fusion using a fused confidence threshold of 0.5 and requiring at least two detections of each object. The eight deployed mine-sized targets all attain high fused confidence values for both surveys, while the ballasts have lower values. The single false alarm (unknown object) is shown in the lower right plot of Figure 5.*

Figure 5 displays ACD images for two example scenes. The reference images (top row) both contain various seafloor texture and objects, as well as schools of fish (one small school in the left scene and four large schools in the right). In the repeat-pass images (middle row) recorded three days later, two targets have been deployed in the left scene. One of the targets is located outside the effective sonar range, in an image region with significant multipath signal pollution. In both change images (bottom row) stationary clutter has been suppressed, thereby enhancing the temporal image differences. The repeat-pass target responses have been preserved (including their trailing highlights reflections from the sea surface), while the fish school highlights in the reference images have been inverted into dark responses. The change images also contain some residual texture, but the matched filter only detects the three mine-sized responses marked with yellow squares. The detection in the right image corresponds to the unknown object in Table 1 and is probably not due to fish, as it looked quite similar in a second view from the same survey. A plausible explanation is an abandoned fishing net fastened in the seafloor concrete blocks, but rearranged by sea currents between surveys 1 and 2. These examples illustrate the effectiveness of ACD for false alarm reduction. Applying the same matched filter directly on the left and right repeat-pass images in Figure 5 produced 12 (including the two targets) and 10 detections, respectively.

Analysis of all single-view detections before fusion reveals four additional detections of the deployed objects (compared to Table 1), distributed between Bottle1-2 in survey 1 and Ballast1 in survey 2. There were also 53 and 187 single-view false alarms (detections not corresponding to any of the 12 deployed objects) for surveys 1 and 2, respectively. This includes a few detections caused by specular reflections from the two recovery ropes. All these additional detections were inconsistent across survey lines and thus removed in the multi-view fusion stage due to missing detections and low confidence values (except the two detections of unknown object in Table 1). The majority of the false alarms appeared to be caused by compact schools of fish (or potentially a single, large fish) of mine-compatible dimensions and often located close enough to the seabed to cast a shadow just behind the sonar highlight (see Figure 6).

The reason for the significantly larger number of false alarms in survey 2 is two-fold. Firstly, the total sonar area coverage is approximately 2.5 times larger than for survey 1. Secondly, the density of near-bottom fish appeared to be higher on the second survey. This interesting result may be due to intra-day variations in fish behaviour (e.g. feeding habits), as the first repeat survey was performed in the afternoon and the second repeat survey mostly before noon. It is difficult, however, to draw firm conclusions on these limited data sets with incompatible survey tracks. For both surveys, the 50 detections with highest single-view confidence included only four non-fish false alarms. These were mostly due to baseline decorrelation and/or co-registration errors in the presence of significant track deviations and rough bathymetry (track offsets are actually visible in the left plot of Figure 4).

Table 2 presents image matching results for the SIFT, SURF, U-SURF (rigid feature orientation) and STAR+BRIEF algorithms applied on the two example scenes in Figure 5. The number of extracted feature points varies between images, but is roughly the same for all methods. Results vary between the scenes, but the number of inlier matches is largest for STAR+BRIEF on both scenes. U-SURF also performs well and SIFT has the highest inliers-to-matches ratio. Similar behaviours have been observed on a larger and more varied SAS data set. SIFT and SURF are designed to tolerate large scene rotations, which is actually a disadvantage in this case, as differences in sonar image orientation are small between repeated surveys.

Algorithm	# Reference feature points		# Repeat pass feature points		# Matches		# Inlier matches	
	Left	Right	Left	Right	Left	Right	Left	Right
SIFT	1810	1086	1543	1102	112	204	95	188
SURF	1742	1288	1528	1339	125	242	61	188
Upright-SURF	1742	1288	1528	1339	157	288	83	245
STAR+BRIEF	1861	1189	1532	1162	246	281	187	249

Table 2: *Results for four feature point algorithms applied on the left and right scenes of Figure 5.*

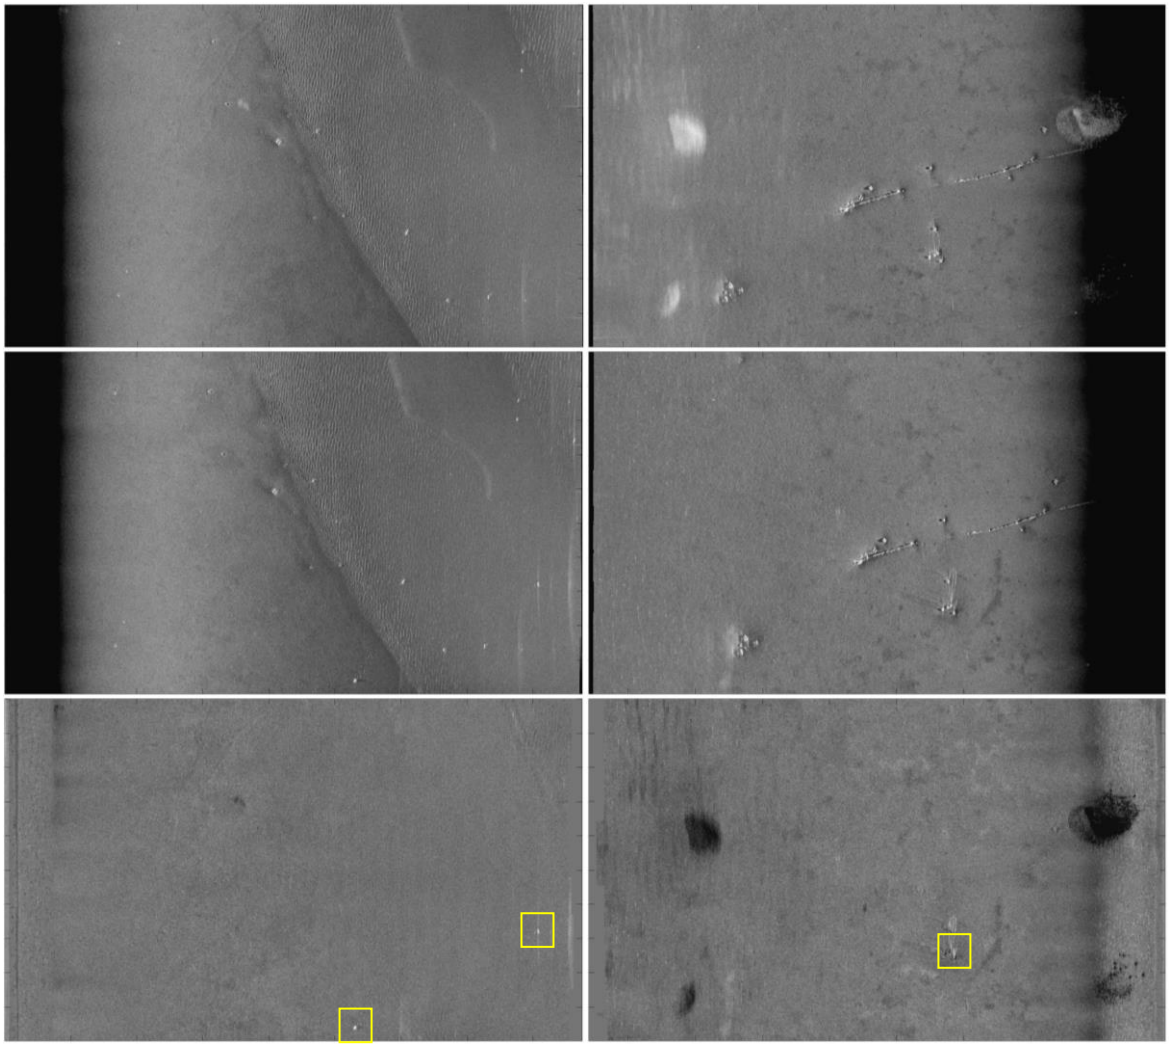

Figure 5: *ACD results for SAS images (175m across-track x 100m along-track) of two scenes containing moderate seafloor roughness and various objects (left column) and concrete blocks, a partly immersed pipeline and fish schools (right column). The plots show the reference-pass images (top row), repeat-pass images (middle) and change images with highlighted detections (bottom). The two detections in the left image correspond to deployed targets MD2 and Rock1, while the detection in the right image is a false alarm, presumably caused by an abandoned fishing net.*

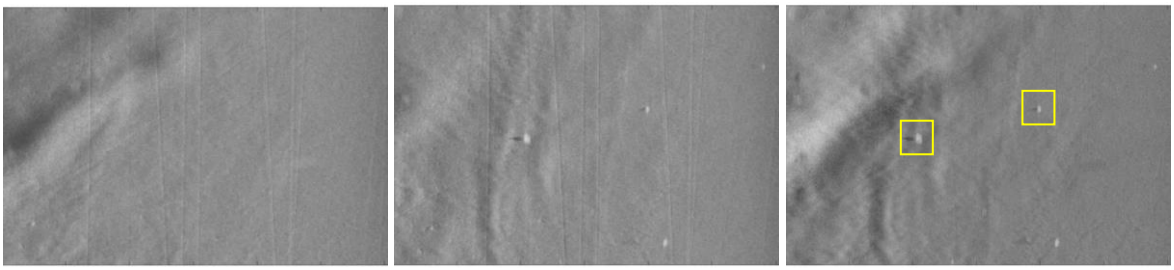

Figure 6: *ACD results for SAS images (40-160m x 75m) showing reference image (left), repeat-pass image (centre) and change image (right) with two highlighted false alarms probably caused by compact schools of fish close to the seabed. Their highlight sizes and trailing shadows resemble target responses. Also visible are two undetected small schools of fish located higher above the seabed (seen by their larger highlight-shadow separation) and several internal wave patterns.*

5 SUMMARY AND FUTURE WORK

This paper presents incoherent automated change detection (ACD) results for HISAS 1030 sonar data recorded with HUGIN-HUS AUV during the MANEX'14 sea trials. The focused application was naval mine hunting and the processing successfully detected all eight deployed, mine-sized objects in two repeat missions with reference time intervals of one and three days, respectively. Target-like fish responses caused multiple false alarms on individual survey lines. These and other spurious false alarms were effectively filtered out by multi-view detection fusion. The processing was performed on streaming SAS imagery, which is the standard HISAS survey product. SAS image generation was not synchronized between the three surveys.

The demonstration of ACD on standard SAS data from a full survey constitutes an important step towards an operational system. The next step will be to test the processing on more full-survey data from varied operations, including surveys with considerably longer time separations. Simultaneously, algorithm development will continue with particular focus on reliable data co-registration in challenging conditions (large track offsets, temporal decorrelation and/or rough bathymetry) and robust single-line detection of new seafloor objects, i.e. improved discrimination of natural scene changes such as fish. As a required module in an operational change detection system for mine hunting, performance evaluation based on *in situ* sensor data will also be developed.

6 ACKNOWLEDGEMENT

The author would like to thank the CMRE personnel and the crew of *NRV Alliance* for all their support during the MANEX'14 exercise.

7 REFERENCES

1. M. L. Gendron and M. C. Lohrenz: "The automated change detection and classification real-time (ACDC-RT) system", in *Proc. MTS/IEEE Oceans Europe*, Aberdeen, UK, June 2007.
2. E. Coiras, J. Groen, D. P. Williams*, et al.*: "Automatic change detection for the monitoring of cluttered underwater areas", in *Proc. of 1st Int. Conf. & Exibition on Waterside Security (WSS)*, Copenhagen, Denmark, August 2008.
3. V. Myers, A. Fortin and P. Simard: "An automated method for change detection in areas of high clutter density using sonar imagery", in *Proc. of Underwater Acoustic Measurements (UAM)*, Nafplion, Greece, June 2009.
4. Ø. Midtgaard, R. E. Hansen, T. O. Sæbø*, et al.*: "Change Detection Using Synthetic Aperture Sonar: Preliminary Results from the Larvik Trial", in *Proc. MTS/IEEE Oceans*, Kona, HI, USA, October 2011.
5. T. O. Sæbø, R. E. Hansen, H. J. Callow*, et al.*: "Coregistration of synthetic aperture sonar images from repeated passes", in *Proc. Underwater Acoustic Measurements 2011*, Kos, Greece, June 2011.
6. V. L. Myers, D. D. Sternlicht, A. P. Lyons*, et al.*: "Automated seabed change detection using synthetic aperture sonar: current and future directions", in *Proc. of the Institute of Acoustics*, vol. 36, pt. 1,pp. 1-10, Lerici, Italy, September 2014.
7. T. G-Michael, B. Marchand, J. D. Tucker*, et al.* "Image-Based Automated Change Detection for Synthetic Aperture Sonar by Multistage Coregistration and Canonical Correlation Analysis", IEEE Journal of Oceanic Engineering, **41**(3): 592-612, July 2016.
8. A. P. Lyons and D. C. Brown. "The impact of the temporal variability of seafloor roughness on the synthetic aperture sonar repeat-pass interferometry", IEEE Journal of Oceanic Engineering, **38**(1): 91-97, 2013.
9. J. Abiva, T. G-Michael, Ø. Midtgaard*, et al.*: "Approaches to false alarm reduction in synthetic aperture sonar change detection", in *Proc. of the Institute of Acoustics*, vol. 40, pt. 2, Lerici, Italy, September 2018.

10. Ø. Midtgaard: "Change detection in synthetic aperture sonar imagery with variable time intervals", in *Proc. 1st Underwater Acoustics Conference and Exhibition*, Corfu, Greece, June 2013.
11. R. E. Hansen, T. O. Sæbø, H. J. Callow*, et al.*: "Synthetic aperture sonar processing for the HUGIN AUV", in *Proc. MTS/IEEE Oceans Europe '05*, Brest, France, June 2005.
12. P. Perona and J. Malik. "Scale-space and Edge Detection Using Anisotropic Diffusion", IEEE Trans. on Pattern Analysis and Machine Intelligence, **12**(7): 629-639, July 1990.
13. H. Bay, A. Ess, T. Tuytelaars*, et al.* "SURF: Speeded up robust features", Computer Vision and Image Understanding (CVIU), **110**(3): 346-359, 2008.
14. D. G. Lowe. "Distinctive image features from scale-invariant keypoints", International Journal of Computer Vision, **60**(2): 91-110, January 2004.
15. M. Agrawal, K. Konolige and M. R. Blas: "CenSurE: Center surround extremas for realtime feature detection and matching", in *Proc. ECCV*,pp. 102-115, Marseille, France, October 2008.
16. M. Calonder, V. Lepetit, C. Strecha*, et al.*: "BRIEF: binary robust independent elementary features", in *Proc. ECCV*,pp. 778-792, Heraklion, Greece, September 2010.
17. G. J. Dobeck: "Fusing sonar images for mine detection and classification", in *Proc. SPIE '99*, vol. 3710,pp. 602-614, Orlando, FL, USA, April 1999.
18. G. Shafer: "A mathematical theory of evidence". Princeton, NJ, USA, Princeton University Press, 1976.
19. R. E. Hansen, A. P. Lyons, T. O. Sæb*ø, et al.* "The effect of internal wave-related features on synthetic aperture sonar", IEEE Journal of Oceanic Engineering, **40**(3): 621-631, July 2015.