CHALLENGES OF AUTOMATED CHANGE DETECTION IN REPEAT-PASS SYNTHETIC APERTURE SONAR IMAGERY

Ø Midtgaard Norwegian Defence Research Establishment (FFI), Kjeller, Norway
T O Sæbø Norwegian Defence Research Establishment (FFI), Kjeller, Norway
N D Warakagoda Norwegian Defence Research Establishment (FFI), Kjeller, Norway

1 INTRODUCTION

Change detection extracts regions of interest as temporal differences by imaging a scene before and after changes may have occurred. We denote the initial image as the reference and the more recent as the current image. Automated tools for image comparison are a prerequisite for many change detection applications, because the large data rates and detailed, yet repetitive contents of the imagery make full manual analysis tedious and prone to errors. Although automated change detection (ACD) in repeat-pass imagery is a well-established technique for synthetic aperture radar (SAR) [1], it is far less mature for underwater sonar due to inherent properties of both traditional side-scan sonar (SSS) imaging and the complex ocean environment. These properties impede for example sensor trajectory control, data positioning accuracy and image resolution. The increasing use of autonomous underwater vehicles (AUVs) equipped with advanced synthetic aperture sonar (SAS) and aided inertial navigation systems (AINS) has partly remedied these limitations and facilitated the development of ACD methods for the sonar domain [2-5]. Still, important challenges remain, in particular regarding scene stability, sensing consistency and image alignment between surveys.

1.1 Applications

ACD in repeat-pass SAS imagery holds the potential to increase significantly both the efficiency and effectiveness of many military and civilian operations connected to seafloor monitoring. The prime naval application is arguably mine counter-measures (MCM), where ACD can improve mine hunting performance by eliminating mutual clutter objects (rocks and debris) in the two datasets. ACD can thus reduce the fraction of seafloor areas currently deemed unhuntable due to excessive clutter densities, and may be the only feasible option to detect mines with unknown physical characteristics, i.e. improvised explosive devices (IEDs) and novel mine models. Other important naval applications include intelligence, surveillance and reconnaissance (ISR), port protection and seabed warfare.

In the offshore industry, ACD can provide operational benefits during external inspection of seafloor infrastructure such as pipelines, cables, well-head and wind farm installations. Another promising application is environmental monitoring of e.g. underwater dumpsites (unexploded ordnance and industrial waste), degradation of hazardous wrecks and trawling/fishing impact on the seafloor. Within marine science, ACD can document e.g. bioturbation activities and vegetation coverage changes.

Due to the diversity of potential ACD applications, the relevant types of changes depend on the given task. Mine hunting, for instance, aims to detect new, mine-like objects, while detections of changes in seafloor texture constitute false alarms. The latter change signatures may, however, be relevant detections for other ACD usages within e.g. environmental monitoring or marine science. To maximize detection performance, it may be beneficial to optimize the ACD algorithm in a supervised manner for the specific task. For many applications, realistic resurvey intervals are in the order of years, due to the high costs of large area seafloor surveys.

The remainder of this paper is structured as follows: Chapter 2 categorizes the ACD processing approaches, while Chapter 3 describes the challenges for ACD in SAS imagery. Chapter 4 visualizes some of these challenges using experimental data from the HISAS1030 sonar. Chapter 5 summarizes the main findings and concludes the paper.

2 PROCESSING APPROACHES

Table 1 categorizes ACD approaches based on the data level used as input for temporal comparison. Both the detection sensitivity and the requirements for successful operation increase considerably from the top to the bottom table entry. Using the most sensitive approach whose requirements are fulfilled in the given environment and scenario, will provide the best achievable performance. Neither of the approaches are therefore generally superior to the others.

Decision level methods process the reference and current data separately to detect and classify e.g. mine-like objects, before matching the two sets of results. The results may include estimated object class, position, dimensions and other characteristics. These approaches are necessarily supervised, because creation of the detector/classifier requires *a priori* models or training data for the given task. Decision level methods are vulnerable to processing errors in the individual data sets and incorrect data associations, but are otherwise unaffected by varying sensing conditions and scene stability. Early sonar ACD works focused on decision level (contact based) methods for mine hunting, i.e. locating new mine-like contacts without a reference contact in corresponding position [6-7].

Contrarily, image level ACD methods compare individual pixels (or small regions) in the current and reference images. These methods can be either non-coherent or coherent, depending on whether they consider only the pixel magnitude or both magnitude and phase. Non-coherent change detection (NCD) is more operationally robust, while coherent change detection (CCD) is more sensitive and may even detect phase changes caused by minute displacements of sediment grains that are undiscernible in the magnitude images. The processing for both variants typically consists of three main steps. The first step retrieves roughly corresponding image data based on the platform navigation solutions and performs algorithm-specific image preprocessing. The next step applies fine-scale image alignment using e.g. cross-correlation [4,5,8] or key point matching [4,9,10] to estimate local or global parameters for image warping. A multi-stage approach may be employed to achieve the required co-registration accuracy, particularly for coherent methods [4,5,8]. The final step extracts changes of interest, typically from a change image where temporal differences have been enhanced. CCD methods use the repeat-pass coherence image, while NCD methods often use the pixel-wise difference (log ratio) image. Dedicated image filtering such as anomaly detectors and automatic target recognition (ATR) algorithms [9] can be applied in order to reduce the number of false alarms from irrelevant image changes. Image level methods can thus be either supervised or unsupervised.

Table 1. *General approaches for automated change detection in repeat-pass SAS imagery. The operational requirements encompass scene stability, sensing geometry/system consistency, image quality and image alignment accuracy.*

Data level	Description	Detection sensitivity	Operational requirements
Decision, Object	Detect differences between corresponding objects in new and old data sets	Medium	Relaxed
Image pixels, Non-coherent	Detect differences in magnitude between corresponding new and old image pixels	Good	Medium
Image pixels, Coherent	Detect differences in magnitude and phase between corresponding new and old image pixels	Excellent	Severe

3 CHALLENGES

Successful realization of operative image-level ACD systems requires addressing several important challenges, which can be divided into four main categories: scene stability, sensing consistency, image degradation and image alignment accuracy. Unresolved, each of these challenges can introduce unwanted differences in repeat-pass image pairs. These irrelevant changes can significantly degrade the detection performance by causing false alarms and masking relevant changes. It is the cumulative effect of the factors that limits ACD performance, such that a reduced sensing consistency (due to e.g. cross-track offset) may demand increased scene stability by shortening the time interval between surveys. All four challenges impose more severe restrictions for CCD than for NCD. In addition, some particular challenges need consideration during ACD system development and implementation.

3.1 Scene stability

The degree of scene stability fundamentally affects the similarity between the reference and current image backgrounds. This factor depends strongly on the ocean environment and seabed properties, but is also a function of sonar frequency and resolution. The degradation of scene stability is often referred to as temporal decorrelation [3]. It is determined by several physical processes that introduce disturbances on the seafloor and in the sound propagation path through the water column [11].

Biological processes reduce the scene stability through bioturbation (e.g. locomotion, feeding and burrowing) and the mere presence of volatile marine organisms such as benthic animals, vegetation and fish. These are persistent, incremental processes, although they typically exhibit seasonal or intraday variations.
Hydrodynamic processes include movement of seabed sediments by bottom currents, waves and turbulence, as well as sound refraction by internal waves in the water column. These may be frequent, incremental processes or rarer, dramatic occurrences, such as storms in shallow waters. Bottom currents may also stimulate dislocation or sedimentary burial of recent seafloor objects.
Anthropogenic processes include human activities like trawling, dredging, anchoring and seafloor infrastructure construction. These processes can radically alter the seabed and usually occur as discrete events. Additionally, wreckage, debris, fish pots and dumped litter constitute sources of new seafloor objects.
Geological processes are less common, but may completely reshape the seabed if they occur. This may be earthquakes or volcanism such as magma eruption and hydrothermal vents.

The scene stability limits the maximum resurvey time interval. For CCD, which relies on the preservation of image speckle, the temporal decorrelation alone may impose extremely stringent restrictions. Recent studies [12] indicate that for a shallow-water sandy seafloor the CCD maximum interval between surveys may be as short as one day, for sonar frequencies above 30 kHz. Many ACD applications have realistic operational resurvey intervals of one year or more. The resulting temporal decorrelation then favors the more robust NCD approaches.

3.2 Sensing consistency

Sensing consistency refers to the similarity of both the sensing geometries and the sonar systems used during the reference and current data collections. Contrarily to the inherently limited scene stability, maintaining the sensing consistency is primarily an operational challenge. Using identical sonar systems to collect the reference and current imagery will provide the most similar images, which again will ensure the most robust ACD results. Any changes in sonar frequency, beam pattern, image resolution, radiometric calibration, image gridding or SAS blocking [9] will affect the sonar images to some degree. For survey intervals of multiple years, software or hardware system variations may be inevitable due to periodic updates and replacement of the survey systems.

Inconsistent sensing geometries are due to horizontal, vertical or orientation offsets between the sensor platform motions in the two surveys. The offsets influence the acoustic response of the seafloor and objects upon it. For instance, the length of a shadow cast by a protruding object depends on both the relative sonar altitude and distance. These image differences are also referred to as baseline decorrelation [1] and can be minimized by using identical mission plans for the surveys. The ability to reproduce the exact reference trajectory is limited, however, by the actual positioning accuracy, stability and manoeuvrability of the platform. Significant topographic variations along the survey path make this especially challenging. Further, variations in sea currents between the surveys may result in different vehicle crab angles, which affect the acoustic response of seafloor objects, including shadow orientation in the images.

3.3 Image degradation

Maintaining high image quality is a major challenge in SAS seafloor imaging [13]. Typical image degradations are defocus and grating lobes, caused by inaccurate estimates of the sonar position, sound velocity or bathymetry. SAS imaging requires a relative positioning accuracy within a fraction of the sonar wavelength along the synthetic aperture length, which typically necessitates navigation techniques correlating sonar data from succeeding pings [14]. This accuracy is particularly demanding to obtain in the case of large topographic gradients or non-linear sensor trajectories. In addition to altering the image intensities, data positioning errors also distort the image gridding.

Acoustical or electronic noise, either from other systems on the sonar platform or from external platforms, can introduce distinct, high intensity patterns in the images or increase the average noise level over large image regions. The images can also be polluted by surface multi-path, typically in shallow waters, reducing the effective signal-to-noise ratio (SNR).

Artefacts like defocus, grating lobes and interference noise usually differ in the two repeat-pass images and thus introduce unwanted changes that degrade the CCD and NCD performances. Even for similar degradations, e.g. multipath pollution in the case of identical tracks and sea states for the two surveys, the resultant reduction in image quality impairs the detection of relevant changes. A particular challenge for CCD, is that image regions with low SNR (sonar shadows, multipath) yield low repeat-pass coherence even without actual changes in the pixel magnitudes.

3.4 Image alignment

Residual misalignments after co-registration and mapping of the repeat-pass imagery materialize as apparent changes in the change image, which may lower the ACD performance. The accuracy requirements are approximately one pixel and one eighth [8] of a pixel for NCD and CCD, respectively. The most common methods for image-based co-registration are cross-correlation (magnitude or complex data) and key point (feature) matching. The methods rely on the existence of persistent scatterers or texture in the image backgrounds. The achievable accuracy thus depends on the scene contents and the degree of change introduced by the above three factors: scene stability, sensing consistency and image degradation. More changes in the backgrounds make image-based co-registration more challenging, until it eventually fails. Again, the processing of complex images are more vulnerable than of magnitude images. Complex image correlation can succeed, however, even on apparently featureless seafloors, provided the image speckle is preserved between surveys.

Even if the data co-registration algorithm succeeds, the image alignment accuracy can be limited by the selected type of image mapping function. While linear transforms (e.g. affine) of the full image are more robust as they average inaccuracies in the individual offset estimates through-out the image, more complicated warping, such as polynomial/splines or local mapping functions, can be required to compensate for non-linear image gridding, especially for CCD approaches. Aligning image pairs with non-linear track differences and/or rough topography can be particularly challenging.

3.5 Development considerations

Development of a supervised ACD system may require a large amount of annotated data to train robust detector algorithms for the given application. Collecting real sonar data from repeat-pass surveys with ground truth of relevant changes, demands considerable resources. Simulated data may be a cost-effective alternative, depending of the number of algorithm parameters to be trained.

Obtaining the ACD results while the current survey is still progressing will provide important operational benefits, such as the ability to immediately inspect the detected changes. Due to the narrow bandwidths of underwater acoustical links, this implies that the ACD algorithms must be sufficiently fast to operate in delayed real-time on an energy-efficient processor in the AUV.

Because the performance of a given ACD system will depend on the environment and operational scenario, establishing a reliable evaluation system for the *in situ* detection performance may be essential for some applications, including mine hunting [15]. The system should take into account the combined characteristics of the environment, the relevant change signatures and the sensor.

4 EXPERIMENTAL RESULTS

We illustrate some ACD challenges using data from various surveys with the HISAS1030 sonar [16] mounted on HUGIN AUVs. The sonar center frequency was 100 kHz. FFI's FOCUS software processed each sonar data set independently (no block synchronization between surveys) into strip map SAS imagery with approximately 4 cm x 4 cm grid cell size. Due to the resurvey interval lengths, we selected an NCD processing chain [9], which co-registered the SAS images by key point matching after speckle filtering and applied an affine mapping of the reference image to produce difference images. All the image presentations apply 50 dB gray-scale dynamics.

Figure 1 shows corresponding port side SAS image sections (120 m x 70 m) from three repeated surveys in a bay outside Horten, Norway. The first two surveys were performed on the 6th and 8th of January 2009 and the third on 14th of January 2019. The resultant resurvey intervals are thus in the extreme limits of two days versus ten years. The scene covers a mud seafloor at around 70 m water depth with a nine meters long wreck in the lower right corner and various scattered debris and rocks. Most of the other image texture is probably caused by bioturbation (langust burrowing, etc.). The images reveal that the objects are stationary over the ten years interval, while the small-scale texture is similar with two days separation, but unrecognizable after ten years.

There are several noticeable issues with the images. The cross track axes reveal a six meters position offset between Figure 1a and the other two images. The sonar data from the first two surveys was regularly polluted by interference noise from the acoustical communication link between the AUV and mother ship. This is evident in the presented imagery as a bright haze in the upper left quadrant of Figure 1b. The tallest part of the wreck is severely defocused in Figure 1c. Finally, some minor, vertical streak artifacts are distinctive to each of the images.

Figure 2 shows the three difference images from the three possible ACD processing combinations of the images in Figure 1. Figure 2a, with only two days separation, reveals the interference pollution and some vertical streaks. In addition, the six meters cross track offset causes baseline decorrelation of the wreck and the debris object at closest range. The background texture in the original images has been well suppressed, except some residual object responses in the lower part of the images. This is probably due to local defocus and/or distorted image gridding. As the interference noise partly obscured the object at x=115 m, y=120 m, this object remains visible in the difference image as a weaker, inversed response, while the two objects above and below it have been suppressed. Similar residual object responses (particularly shadows) would occur for the case of two repeat-pass images with different levels of multipath pollution.

The two difference images with ten years separation are quite similar to each other, except for the baseline decorrelation of the wreck and close range debris object in Figure 2b and the interference pollution (which is now dark because it is in the reference image) in Figure 2c. The images are dominated by the background sediment texture, which has changed radically over ten years. Also, details on the wreck are visible as changes, due to the blur in Figure 1c. A striking feature of these two difference images is the large debris object visible near the upper left corner. Examining the original survey images, we notice that the object has somehow changed its orientation between the second and third survey and thus correctly appears as a change in the difference images. Due to its signature uniqueness, this object change is clearly distinguishable from all the small-scale texture changes elsewhere in the difference images.

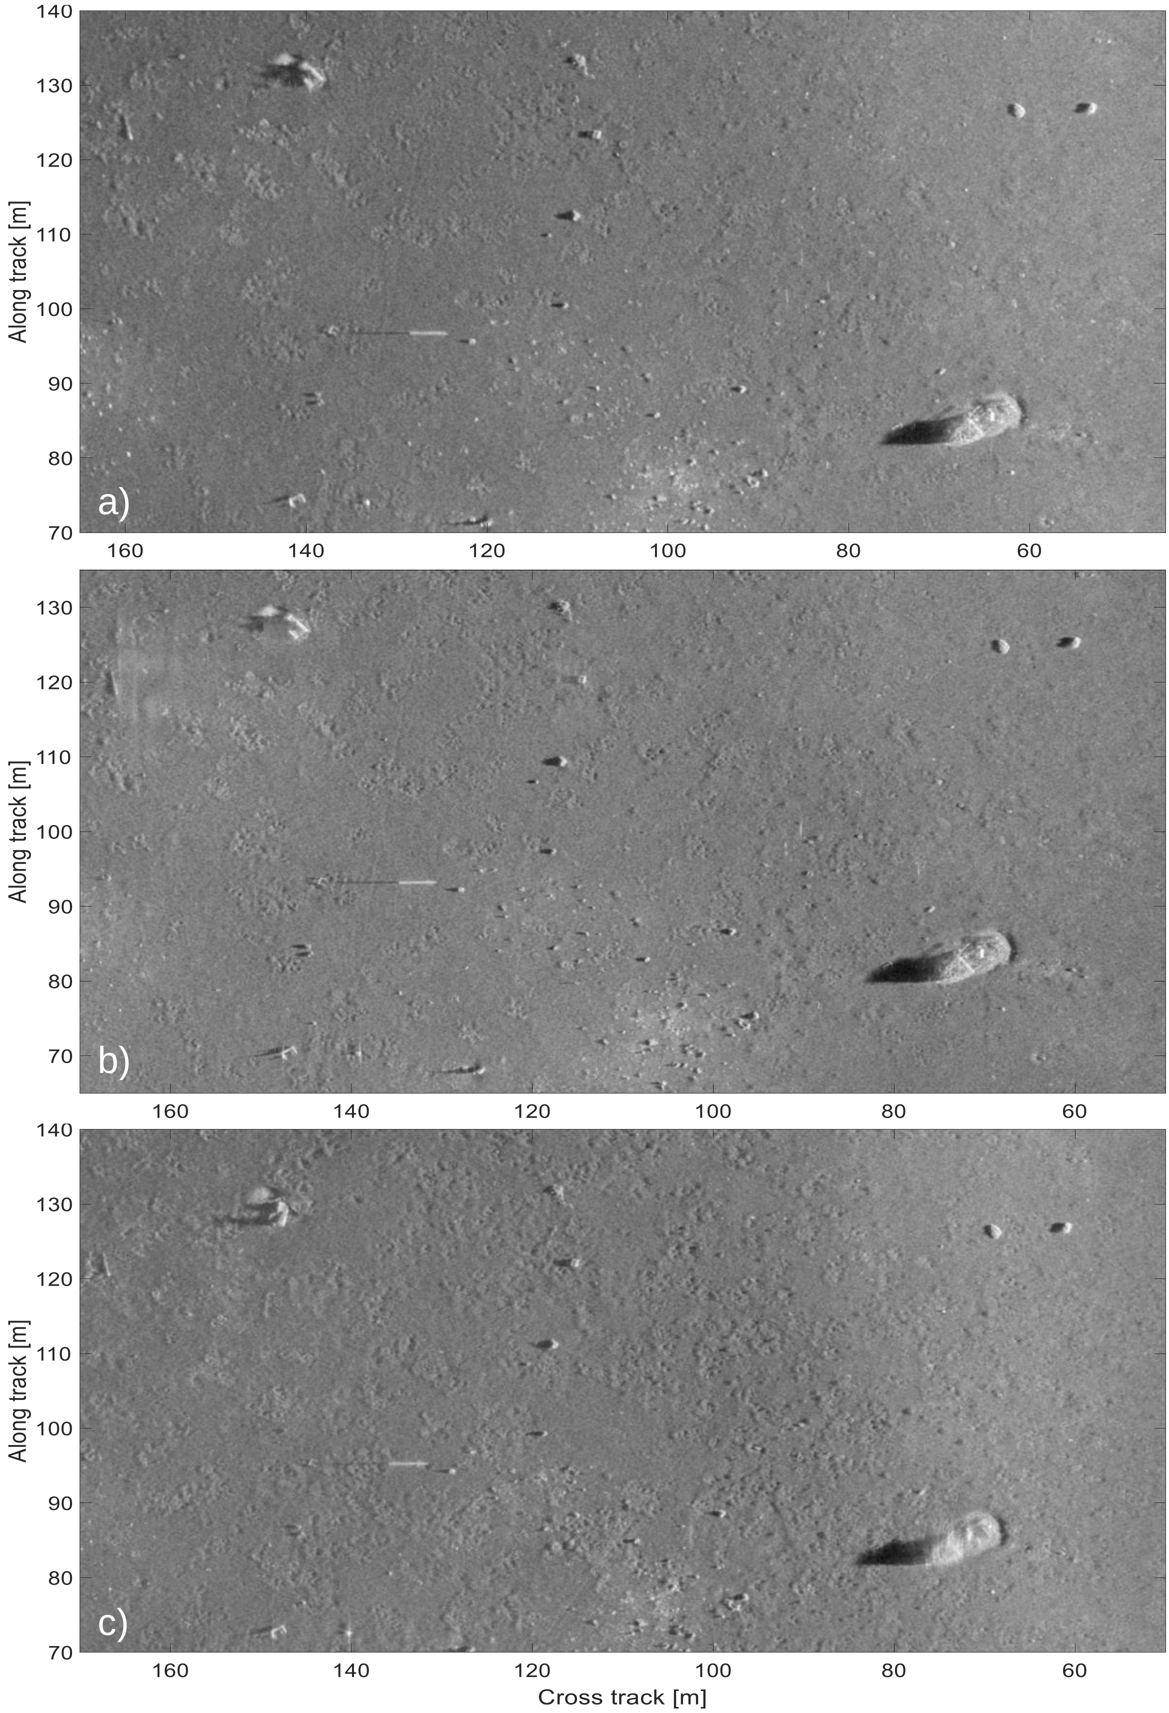

Figure 1. *SAS image sections from survey dates a) 06-Jan-2009, b) 08-Jan-2009, c) 14-Jan-2019.*

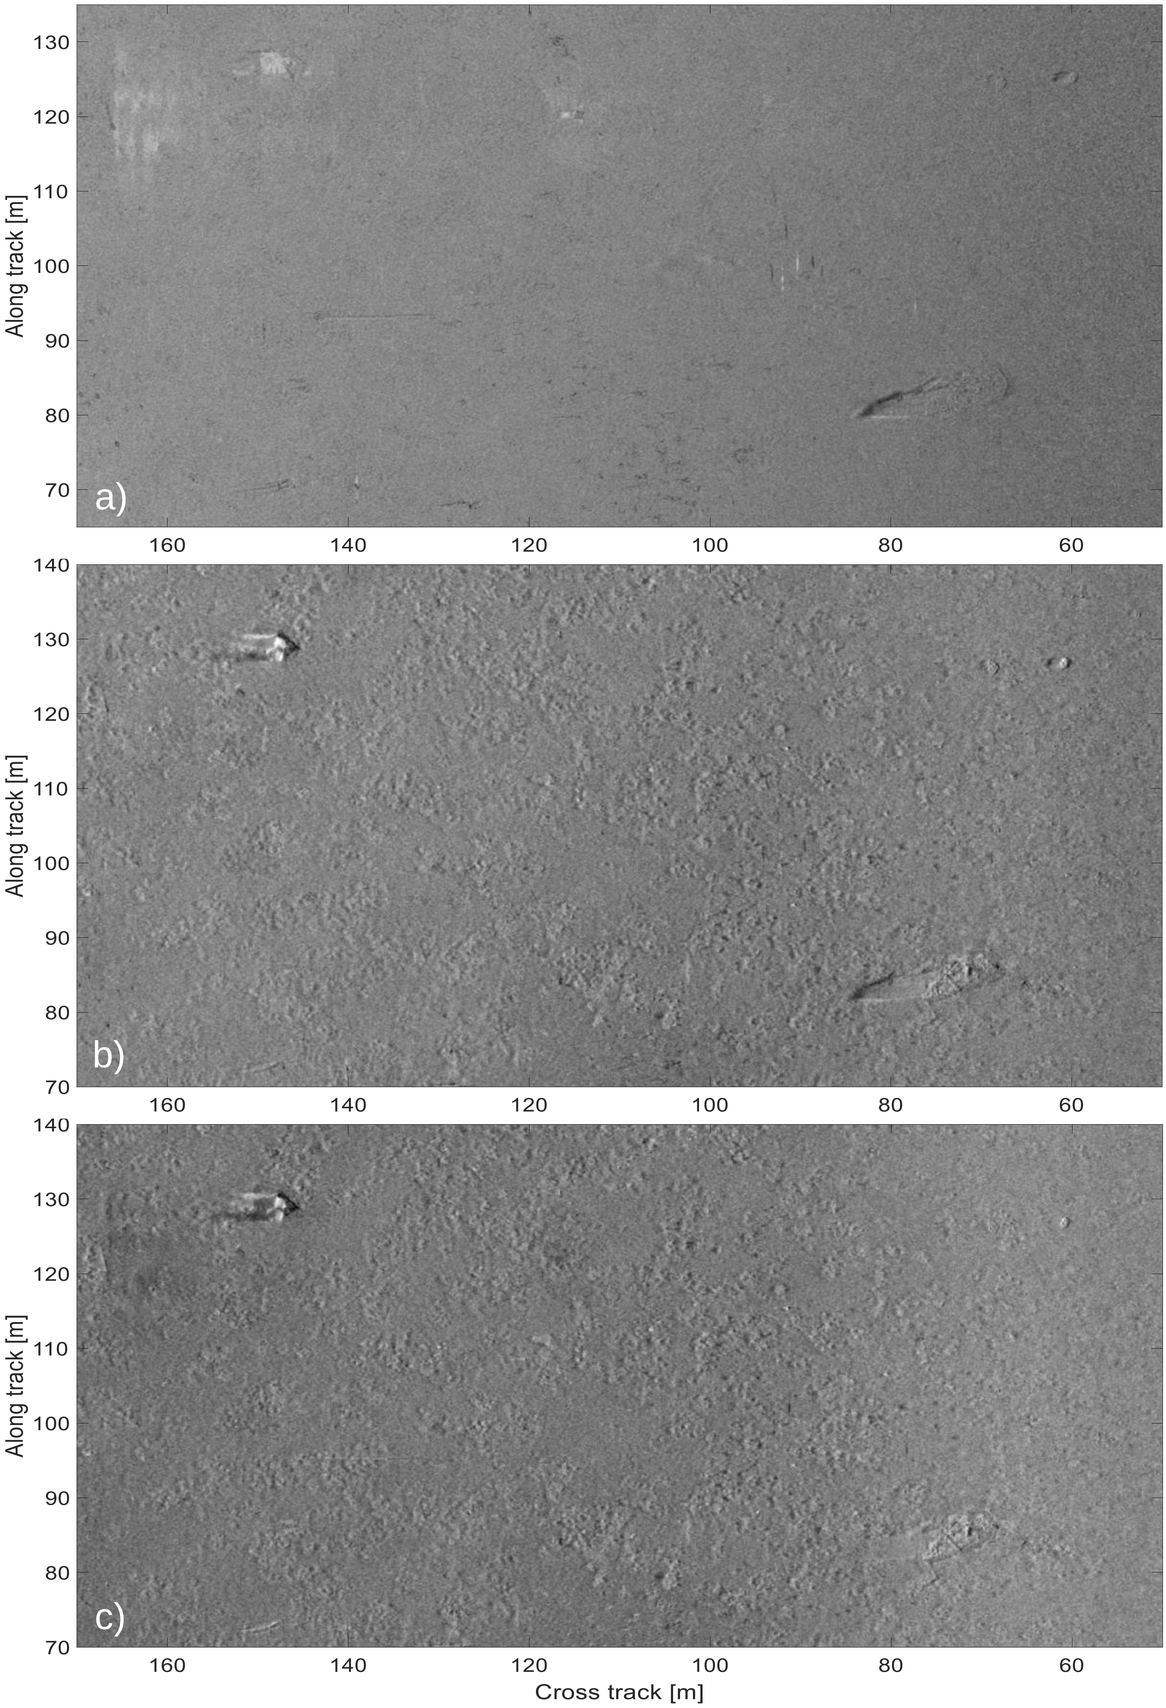

Figure 2. *Difference images of aligned image sections in Figure 1. a) 1b-1a. b) 1c-1a. c) 1c-1b.*

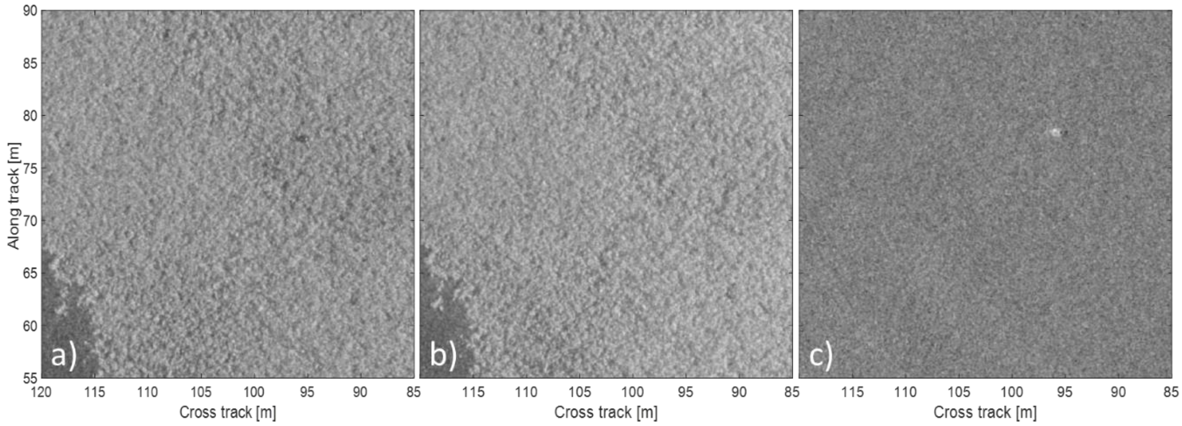

Figure 3. *Port SAS image sections (35 m x 35 m) of a Posidonia seagrass scene before and one year after retrieval of test object (car wheel). a) Reference image with object. b) Repeat-pass image without object. c) Difference image revealing inversed highlight and shadow of the removed object.*

Figure 3 presents HISAS data collected off the island of Elba, Italy, during the ARISE'12 and MANEX'13 sea trials arranged by the NATO Centre for maritime research and experimentation (CMRE). The 35 m x 35 m image sections show a *Posidonia* seagrass meadow with a patch of exposed sand sediments in the lower left corner. The water depth is 25-30 m. The survey dates were 18th October 2012 and 10th October 2013, yielding a time separation of approximately one year. A car wheel (tyre and rim) was deployed on the seafloor before the reference survey and retrieved a few days after. The stability of this vegetated scene with one year survey interval was expected to be challenging for ACD, regarding provision of stationary image texture for co-registration. Results showed, however, that a sufficiently large number of key points were successfully matched in the full scene. Most of these points were located in the transition between seagrass and sediments, but many were also found within the meadow. The seagrass texture have been suppressed in the difference image, thus enhancing the inversed response of the removed wheel.

5 CONCLUSIONS

The operational use of ACD in repeat-pass SAS imagery requires addressing several challenges, which can be categorized as: 1) scene stability between surveys, 2) consistency of sensing geometries and systems, 3) image degradation and 4) image alignment accuracy. Unresolved, these challenges typically introduce irrelevant changes between the aligned repeat-pass images, thereby potentially causing false alarms and masking relevant changes. The challenges apply for all image-level ACD approaches, but the requirements are significantly stricter for CCD than NCD. Due to severe scene stability requirements, CCD methods may not be compatible with the operational realistic resurvey intervals for many ACD applications.

This paper presents experimental results from NCD processing of HISAS1030 images from two challenging scenarios. The first example demonstrated detection of a removed object (car wheel) in a shallow water seagrass meadow with a survey separation of one year. The second example showed that NCD can be feasible on a mud seafloor at 70 m water depth even for the longest tested survey interval of ten years. Generally, the maximum viable survey separation depends strongly on the environment and operational scenario. In addition to challenges from temporal background changes, the second set of repeat-pass imagery contained interference noise, cross-track offset and image defocus. These real data examples demonstrate that NCD methods can tolerate considerable image background changes, provided there is sufficient persistent texture for image-based co-registration and that the signatures of relevant changes can be distinguished from those of irrelevant changes. However, if the total amount of irrelevant changes causes the change images to contain more clutter than the repeat-pass images, it may be favourable to analyse the single-pass images only.

6 REFERENCES

1. M. Preiss and N.J.S. Stacy, “Coherent change detection: theoretical description and experimental results”, DSTO-TR-1851, Australia (2006).
2. Ø. Midtgaard, “Change detection in synthetic aperture sonar imagery with variable time intervals”, Proc. 1st Underwater Acoustics Conf., Corfu, Greece (2013).
3. V.L. Myers, D.D. Sternlicht, A.P. Lyons and R.E. Hansen, “Automated seabed change detection using synthetic aperture sonar: current and future directions”, Proc. Inst. of Acoustics, 36(1), Lerici, Italy (2014).
4. 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 J. Ocean. Eng., 41(3), 592-612 (2016).
5. V. Myers, I. Quidu, B. Zerr, et al., “Synthetic aperture sonar track registration with motion compensation for coherent change detection”, IEEE J. Ocean. Eng., 45(3), 1045-1062 (2020).
6. M.L. Gendron and M.C. Lohrenz. “The automated change detection and classification real-time (ACDC-RT) system”, Proc. MTS/IEEE Oceans Europe Conf., Aberdeen, UK (2007).
7. E. Coiras, J. Groen, D.P. Williams, et al., “Automatic change detection for the monitoring of cluttered underwater areas” Proc. 1st Waterside Security Conf., Copenhagen, Denmark, (2008).
8. T.O. Sæbø, R.E. Hansen, H.J. Callow and S.A. Synnes, “Coregistration of synthetic aperture sonar images from repeated passes”, Proc. Underwater Acoustic Measurements Conf., Kos, Greece (2011).
9. Ø. Midtgaard, “Automated change detection in streaming SAS imagery”, Proc. Inst. of Acoustics, 40(2), Lerici, Italy (2018).
10. V.T. Wang and M.P. Hayes, “Synthetic aperture sonar track registration using SIFT image correspondences”, IEEE J. Ocean. Eng., 42(4), 901-913 (2017).
11. E. Shouldice, S. Smith, T.O. Sæbø, et al., “The impact of scene stability and acoustic backscatter coherence on automated seabed change detection”, Proc. Inst. of Acoustics, 45(1), Lerici, Italy (2023).
12. 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 J. Ocean. Eng., 38(1), 91-97 (2013).
13. R.E. Hansen, H.J. Callow, T.O. Sæbø and S.V. Synnes, “Challenges in seafloor imaging and mapping with synthetic aperture sonar”, IEEE Trans. Geosci. Remote Sens., 49(10), 3677-3687 (2011).
14. A. Bellettini and M. Pinto, “Theoretical accuracy of synthetic aperture sonar micronavigation using a displaced phase-center antenna”, IEEE J. Ocean. Eng., 27(4), 780-789 (2002).
15. D.D. Sternlicht, A. Crawford, T. G-Michael, et al., “Operator tools and performance assessment for automated seabed change detection”, Proc. Inst. of Acoustics, 40(2), Lerici, Italy, (2018).
16. P.E. Hagen, T.G. Fossum and R.E. Hansen, “HISAS 1030: The next generation mine hunting sonar for AUVs”, Proc. Undersea Defence Technology Pacific Conf., Sydney, Australia (2008).