ON THE UTILITY OF COHERENCE FOR UNDERWATER OBJECT DETECTION IN SAS IMAGES

David P. Williams

NATO Science and Technology Organization, Centre for Maritime Research and Experimentation, Viale San Bartolomeo 400, 19126 La Spezia, Italy

Email: williams@cmre.nato.int; fax: +39 0187-527-330

Abstract: *A large-scale study is undertaken to assess the capability of detecting underwater objects in sonar imagery. The data used in the analysis are synthetic aperture sonar (SAS) images containing various man-made objects. The data were collected with the MUSCLE autonomous underwater vehicle (AUV) during six large sea trials, conducted between 2008 and 2012, in different geographical locations with diverse environmental conditions. The detection algorithm for which performance is assessed is a cascaded, integral-image-based approach. The analysis examines detection performance for specific target shapes as a function of target-sensor aspect, range, image quality (based on spatio-temporal coherence), and environment (defined in terms of two seabed features). To our knowledge, this study – covering approximately 158 square kilometers of seabed imagery, and involving over 1400 target detection opportunities – represents the most extensive such systematic, quantitative assessment of object detection performance with SAS data. The analysis reveals that, for detection performance, the image quality (or spatio-temporal coherence) is a more important quantity than range, aspect, or shape. Dependence on the environment, and in particular seabed characterized by sand ripples, is also established. It is then explained how these results can be exploited, by intelligently adapting AUV survey routes, to maximize detection performance over an area of interest.*

Keywords: *Detection, synthetic aperture sonar (SAS), image quality, performance assessment, autonomous underwater vehicle (AUV), adaptive surveying.*

1. INTRODUCTION

Using an autonomous underwater vehicle (AUV) equipped with side-looking sonar – and in particular, high-resolution synthetic aperture sonar (SAS) – has become a feasible approach for detecting objects on the seafloor. However, the subsequent performance assessment associated with this task often relies on subjective, qualitative estimates of environmental characteristics, such as seabed composition, that are notoriously challenging to predict accurately. Alternative approaches that employ theoretical sonar prediction models to estimate performance do not overcome this problem, for the models are extremely sensitive to estimated parameters that are difficult or impossible to measure accurately *in situ*. We argue that performance assessment would be more reliable if it were instead based on rigorous quantities computed using through-the-sensor data. To this end, in this work, we seek to establish the quantities that have the most significant impact on object detection performance. This information could then be used to develop a new approach for performing AUV planning and evaluation with high-resolution SAS imagery. The results can also be exploited to ensure collection of the most valuable data at sea.

The remainder of this paper is organized as follows. Sec. 2 describes the SAS data that are used in the study. Sec. 3 provides an overview of the quantities for which detection performance is assessed. The detection results are shown in Sec. 4, and ways to exploit the findings are given in Sec. 5.

2. DATA

This study examines object-detection performance using sonar data collected during six major sea trials conducted by NURC/CMRE between 2008 and 2012; the trials were Colossus 2 (2008), CATHARSIS 1 (2009), CATHARSIS 2 (2009), AMiCa (2010), ARISE 1 (2011), and ARISE 2 (2012). This collection of data spans diverse environments in terms of seafloor characteristics, including flat hard-packed sand, soft mud, seabed characterized by sand ripples, and seabed covered in posidonia. In each sea trial, different groups of objects (i.e., targets) were laid, and AUV surveys were then performed over the target areas. The principal target classes considered here are cylinders, truncated cones, and wedge-shaped objects; other objects are also included in the analysis. In all, the study spans SAS imagery covering an area of approximately 158 square kilometers of seabed and involves over 1400 target detection opportunities.

All of the data used in this study were collected by the CMRE's SAS-equipped MUSCLE AUV. The center frequency of the SAS is 300 kHz and the bandwidth is approximately 60 kHz. The system allows the formation of high-resolution sonar imagery with a theoretical across-track resolution of 1.5 cm and a theoretical along-track resolution of 2.5 cm.

3. EXPERIMENTAL SET-UP

The target-detection algorithm described in [1] is used as the basis for detection performance assessment. The algorithm is an integral-image-based approach that makes (delayed) real-time detection possible onboard an AUV. The cascaded algorithm architecture is designed to minimize computational costs by operating on progressively

smaller portions of an image at each of three major stages: shadow detection, ripple detection, and echo detection. The algorithm exploits extensive domain-specific knowledge and is tailored to the fundamental underlying physics and geometry of the problem. (In addition, the algorithm can also take image quality into consideration. In regions with poor image quality, where shadows are expected to be impure and hence unreliable, detection can instead be based solely on echoes.)

The detection performance as a function of several quantities is studied. These quantities include the target-sensor aspect, the range from the sensor to the target, the image quality, and two environmental features – the anisotropy and complexity of the seabed.

The mean correlation value of consecutive ping returns, as a function of range, is employed as a surrogate measure of the image quality. This quantity is computed within the displaced phase-center antenna (DCPA) algorithm [2] of the SAS processing and effectively exploits both spatial and temporal coherence properties of the problem; for this reason, we refer to it as spatio-temporal coherence, or simply a generic coherence.

The two environmental features considered measure the anisotropy and complexity of the seabed, which were introduced in [3]; modifications were made to the manner in which the features were computed, with this described in [4]. The anisotropy feature measures the *variation* of filter responses with an image in different directions. For the SAS images, this effectively quantifies the degree to which the seabed exhibits directional features. The complexity feature measures the average of filter responses with an image in different directions. For the SAS images, this effectively quantifies the degree to which the seabed contains objects or other sources of irregularity.

To ensure that the presence of the object on the seabed does not skew the environment features, the responses of the detections are masked prior to the feature calculations. As a result, these environment features should be independent of the object around which they are computed. Evidence of the ability of these features to capture the characteristics of the seabed is provided in Fig. 1. Specifically, we show the distribution of these features in three distinct environments: (i) on flat seabed, (ii) on seabed characterized by sand ripples, and (iii) on the boundary between flat and rippled seabed. In all three cases, the features are extracted from the area of seabed around the same target shape, namely a truncated cone. For each case, multiple views of the object were obtained from various aspects. One example SAS image from each of the three environments, along with the distribution of the features, is also shown in Fig. 1.

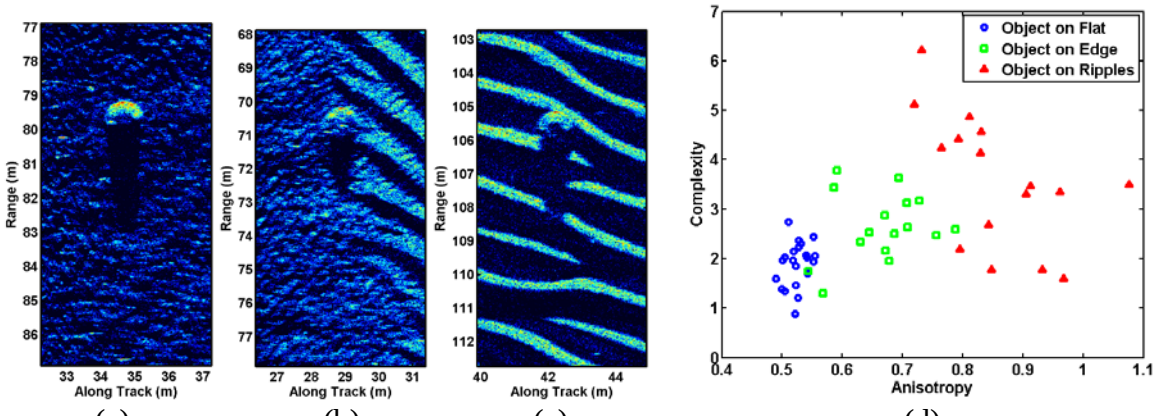

(a) (b) (c) (d)

Fig.1: Example SAS image of the same target shape (a) on flat seabed, (b) on the boundary between flat seabed and ripples, and (c) on seabed characterized by sand ripples; (d) the distribution of the two environment features extracted from the various views of this target shape at these three locations.

4. RESULTS

Detection performance as a function of various quantities for all of the data described in Sec. 2 is presented here. It should be noted that the case indicating "All" in the figures encompasses several additional target types beyond the three specific shapes named. The average number of detection opportunities for each of the three specific shapes is about 300; the total number of detection opportunities for all targets is about 1400. In all of the subsequent figures, the x-axis tick marks and labels indicate the bin ranges used when quantizing the given quantity under study.

Detection performance as a function of image quality, or what we term spatio-temporal coherence, is shown in Fig. 2. It can be seen that detection performance exhibits a strong dependence on image quality, which can be attributed to the fact that shadows become unreliable at lower image-quality values because they tend to "fill in." This insight suggests that image quality may also be a useful parameter to consider in other algorithms that exploit shadows, such as segmentation methods used for feature extraction.

If the detection algorithm takes image quality into account, effectively ignoring shadow information, the detection rates at lower image quality values (below 2/3) can be raised substantially. (The performance reflecting this version is indicated in Fig. 2 with markers.) The remaining results are based on this more effective version of the detection algorithm.

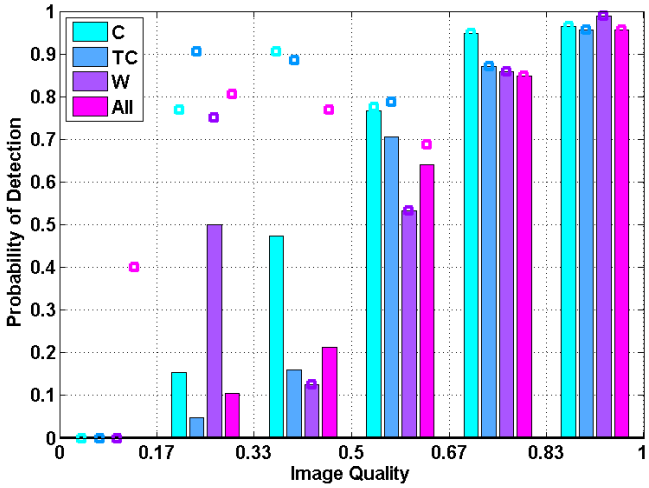

Fig.2: Detection performance for each object type across all trials, as a function of image quality (C=cylinders, TC=truncated cones, W=wedge-shaped targets, All=all targets). The markers indicate performance when the detection algorithm is modified to ignore shadows and rely only on echoes when the image quality is below 2/3.

Detection performance as a function of range is shown in Fig. 3. In general, range does not strongly impact detection performance in SAS imagery, which makes sense because the resolution in SAS images is independent of range. This is in contrast to side-scan sonar imagery, where the resolution – and hence, detection performance – is indeed range-dependent. Therefore, for planning and evaluation purposes, constructing models for detection probability as a function of lateral range – so-called "P(y) curves" – is less helpful when dealing with SAS data (compared to side-scan data). The results in Figs. 2-3 are in agreement with another smaller scale study [5] that also found image quality (based on coherence) to be more relevant than range for detection purposes.

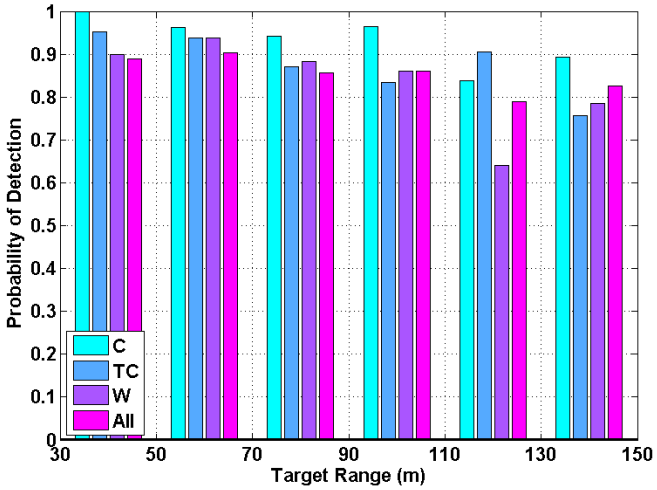

Fig.3: Detection performance for each object type across all trials, as a function of range (C=cylinders, TC=truncated cones, W=wedge-shaped targets, All=all targets).

Fig. 4 shows the detection performance as a function of the two environment features, seabed anisotropy and seabed complexity. Based on the results in Fig. 1, it was clear that the division between non-rippled seabed and rippled seabed occurred near an anisotropy value of 2/3. Therefore, Fig. 4(a) shows the detection performance separated at this threshold. High anisotropy values, which are indicative of sand ripples, corresponded with a lower detection rate. This result makes intuitive sense because the shadows cast by targets can merge with the shadows cast by ripples, making detection difficult. Similarly, the highlight of a target can blend in with the highlight due to ripples, which effectively obscures the targets. In terms of complexity, a low value is indicative of a benign seabed, so it is not surprising that detection rates were high in that regime. Taken together, the results of anisotropy and complexity support the long-held view that detection capability is a strong function of the environment. Being able to reliably establish the environment quantitatively using through-the-sensor data, however, is a key advance.

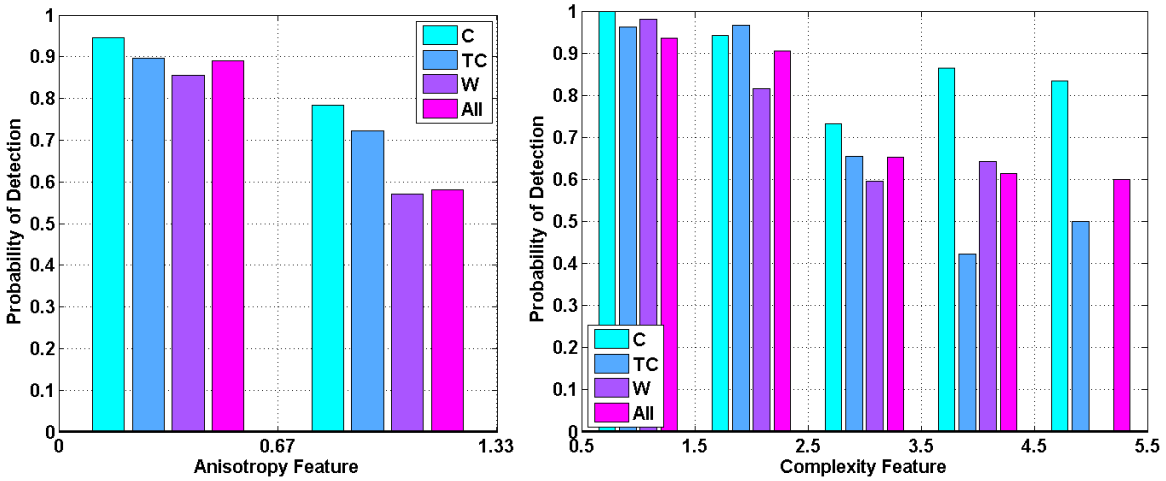

(a) (b)

Fig.4: Detection performance for each object type across all trials, as a function of seabed (a) anisotropy and (b) complexity (C=cylinders, TC=truncated cones, W=wedge-shaped targets, All=all targets).

Although one cannot necessarily change the environment in which the detection is to be performed, additional results suggest that collecting data at certain target-sensor *aspects* can mitigate the adverse effects of the environment. In Figs. 5-7, detection performance is shown as a function of target-sensor aspect on polar-coordinate plots. In these figures, the radius and angle of the plot indicate the probability of detection and the target orientation, respectively. Concentric grid-rings are shown on these plots in 0.25 increments. The space of target orientations is divided into four non-overlapping sets each spanning 90°, corresponding to the front, rear, and each side of the target.

In Fig. 5, we present results for the cylinder and wedge-shaped targets – the truncated cone is rotationally symmetric – considering only instances where the objects were on flat seabed (i.e., excluding instances in ripple fields). For both target shapes, very little aspect dependence was observed. This result is somewhat surprising because it is generally accepted that, for example, cylinders viewed at endfire (where there are fewer pixels on target) are more challenging to detect than cylinders viewed at broadside. Instead, the results suggest that the more important quantity for determining detection performance is the image-quality score.

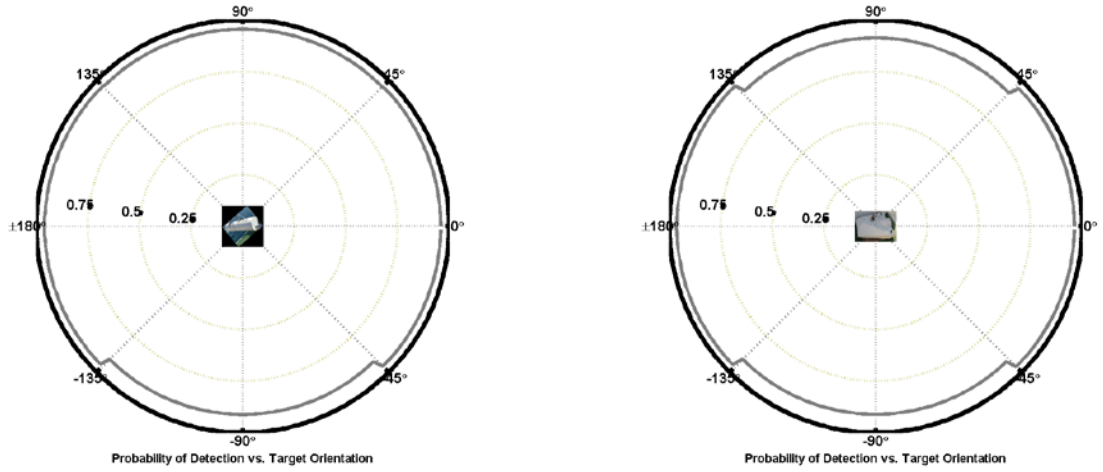

(a) (b)

Fig. 5: Detection performance across all trials, on flat benign seabed, as a function of target orientation, for (a) cylinder targets and (b) wedge-shaped targets.

The more interesting results concerning target aspect are for the cases of targets in ripple fields. It was observed in Fig. 4(a) that detection capability has a strong dependence on the presence of sand ripples. However, the *orientation* at which the sand ripples are surveyed also has a dramatic impact on performance. This can be best observed in Figs. 6-7, which correspond to targets laid in sand ripple fields. (These results are admittedly only anecdotal, based on a small sample-size of about 20 detection opportunities per case.) In Fig. 6, the targets are rotationally symmetric truncated cones, but detection performance is still aspect dependent; specifically, detection performance is much better when the sonar viewing aspect is looking along a ripple crest (or trough). Therefore, when sand ripples are present, it is recommended that one surveys such that the sonar looks along ripple crests (and troughs) rather than across them. In this way, one minimizes the chance that targets

will be obscured by the ripples, either from target and ripple highlights melding or from ripple shadow concealment.

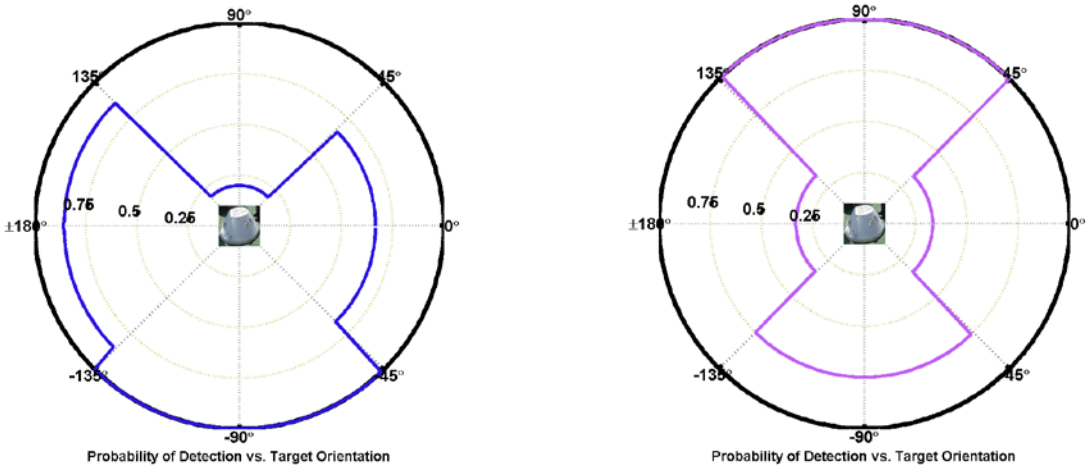

Fig. 6: Detection performance for two truncated cones as a function of target orientation. The targets were laid in sand ripple fields oriented at (a) 160°/-20° and (b) 130°/-50°.

The results in Fig. 7 also reveal another interesting phenomenon regarding detection performance of targets in ripples. That figure shows the results for the same target type (cylinder), with the major difference being only the relative orientations of the targets with respect to the sand ripples. The target associated with the results in Fig. 7(a) was oriented nearly parallel to the ripple crests, whereas the target associated with the results in Fig. 7(b) was oriented nearly orthogonal to the ripple crests. This arrangement produced strong aspect dependence for detection in the former case, but no aspect dependence in the latter case. Collectively, these results suggest that the relative orientation of targets to sand ripples also has a major impact on detection ability, with the most challenging case being when the principal axis of the target and a ripple crest align. This result makes intuitive sense because it is in such cases that the highlights of the target and ripples would coincide, effectively masking the target response.

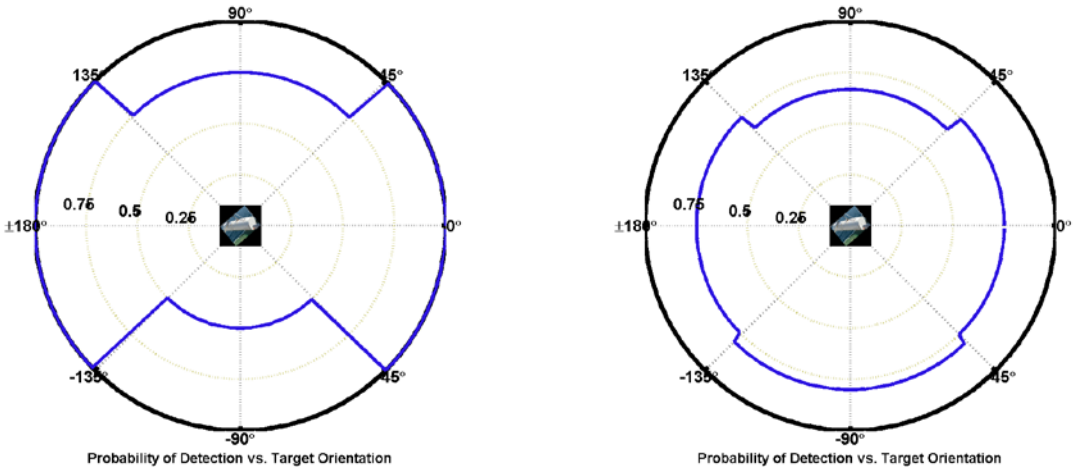

Fig. 7: Detection performance for two cylinders as a function of target orientation. The targets were laid in sand ripple fields oriented at (a) 175°/-5° and (b) 115°/-65°.

5. CONCLUSION

Although this study examined the performance of only one particular detection algorithm, the results indicate that strong dependence with certain quantities does exist. Moreover, the results presented here can now serve as substantial evidence to support the intelligent adaptation of AUV surveys for improving target detection performance.

The most common reasons that targets were not detected were due to poor image quality, concealment by sand ripples, and imaging at unfavorable or "unlucky" target aspects. These issues can be mitigated by intelligent surveying in the following ways. To guard against poor image quality, the tracks of an AUV survey can be adapted to ensure that high image-quality data are collected over the entire area of interest [6, 7]. Regarding sand ripples, once they are detected, the principal orientation of the AUV survey can be adjusted so that the tracks are oriented orthogonal to a given sand ripple crest (i.e., parallel to the direction in which the ripple "waves" propagate) [8]. To prevent viewing targets at only unfavorable aspects, an additional survey that exhibits substantial aspect diversity (such as a 90° orientation difference) can be executed [9].

In addition, the results hint that segmentation algorithms used to extract features for the subsequent classification stage may benefit from using knowledge about the image quality, since the detection analysis obliquely revealed that the purity of shadows (and hence the ease of segmentation) is a strong function of image quality.

REFERENCES

[1] **D. Williams and J. Groen**, "A Fast Physics-Based, Environmentally Adaptive Underwater Object Detection Algorithm," in *Proc. OCEANS*, 2011.

[2] **A. Bellettini and M. Pinto**, "Theoretical Accuracy of Synthetic Aperture Sonar Micronavigation Using a Displaced Phase-Center Antenna," *IEEE Journal of Oceanic Engineering*, Vol. 27, No. 4, pp. 780-789, 2002.

[3] **O. Daniell, Y. Petillot, and S. Reed**, "Unsupervised Sea-Floor Classification for Automatic Target Recognition," in *Proc. International Conference on Underwater Remote Sensing*, 2012.

[4] **E. Fakiris, D. Williams, M. Couillard, and W. Fox**, "Sea-Floor Acoustic Anisotropy and Complexity Assessment Towards Prediction of ATR Performance," in *Proc. International Conference on Underwater Acoustics*, 2013.

[5] **S. Synnes, R. Hansen, and T. Sæbø**, "Assessment of Shallow Water Performance Using Interferometric Sonar Coherence," in *Proc. International Conference on Underwater Acoustic Measurements*, 2009.

[6] **D. Williams**, "On Optimal AUV Track-Spacing for Underwater Mine Detection," in *Proc. IEEE International Conference on Robotics and Automation*, pp. 4755-4762, 2010.

[7] **D. Williams, A. Vermeij, F. Baralli, J. Groen, and W. Fox**, "In Situ AUV Survey Adaptation Using Through-the-Sensor Sonar Data, *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2012.

[8] **D. Williams**, "AUV-Enabled Adaptive Underwater Surveying for Optimal Data Collection," *Intelligent Service Robotics*, Vol. 5, No. 1, pp. 33-54, 2012.

[9] **M. Couillard, J. Fawcett, and M. Davison**, "Optimizing Constrained Search Patterns for Remote Mine-Hunting Vehicles," *IEEE Journal of Oceanic Engineering*, Vol. 37, No. 1, pp. 75-84, 2012.