A Fast Physics-Based, Environmentally Adaptive Underwater Object Detection Algorithm

David P. Williams and Johannes Groen
NATO Undersea Research Centre, Viale San Bartolomeo 400, 19126 La Spezia (SP), Italy
{williams, groen}@nurc.nato.int

***Abstract*—A new algorithm for the detection of underwater man-made objects in sonar imagery is proposed. The algorithm is made extremely fast by employing a cascaded architecture and by exploiting integral-image representations. As a result, the method makes real-time detection of objects of interest in streaming sonar data collected by an autonomous underwater vehicle feasible. No training data is required because the proposed method is adaptively tailored to the environmental characteristics of the sensed data that is collected *in situ*. The flexible yet rigorous approach also addresses and overcomes five major limitations that plague the most popular detection algorithms that are in common use. The power and utility of the proposed approach is demonstrated on a large, challenging data set of synthetic aperture sonar imagery collected at sea.**

***Index Terms*—Detection, underwater objects, synthetic aperture sonar.**

I. Introduction

The high-resolution imaging of underwater environments afforded by sonar has proven particularly useful for the detection of man-made objects on the seabed. Thanks to breakthroughs in marine robot technology, the sonar data used to address this task is invariably collected by an autonomous underwater vehicle (AUV).

Because of the inherent danger and time-sensitive nature of these costly operations, the next urgent priority is to embed intelligence in the AUV so that it can immediately react to the data it collects. By adapting its survey route *in situ* and efficiently allocating resources, the AUV can collect the most informative data for the task at hand while simultaneously reducing costs.

Before this goal can be realized, however, an algorithm is needed that can perform robust object detection in near-real-time onboard an AUV with limited processing capabilities. The existing detection algorithms in widespread use suffer from several limitations that preclude this possibility.

In this work, a novel algorithm for the detection of underwater man-made objects in synthetic aperture sonar (SAS) imagery is presented. The proposed method is an adaptive algorithm that explicitly accounts for environmental characteristics, image quality, and other physics-based aspects of the problem. By incorporating adaptivity into the algorithm, the flexible yet rigorous approach also addresses and overcomes five major limitations that plague the most popular detection algorithms that are in common use.

The remainder of this paper is organized as follows. Sec. II provides background about the underwater object detection task and notes the limitations of existing detection algorithms. The proposed cascaded detection algorithm is described in Sec. III. Experimental results on real, measured sonar data collected at sea are presented in Sec. IV, followed by a brief discussion in Sec. V. Concluding remarks and directions for future work are noted in Sec. VI.

II. Underwater Object Detection

A. Background

The high-resolution imaging of underwater environments afforded by SAS [1] is particularly useful for the detection of man-made objects on the seabed. The need to perform such detection in an automated manner with no human intervention is motivated by the desire to conduct fully autonomous detection missions with an AUV.

Once the sonar data has been collected and processed into imagery, an object detection algorithm is applied. Typically, these detection algorithms search for highlight-shadow patterns characteristic of man-made objects of interest. The highlight is the result of the acoustic echo from the object itself, while the shadow that is cast is due to the geometry between the object (and specifically its height above the seafloor) and the grazing angle of the transmitted signal.

B. Motivation

Existing underwater object detection algorithms [2]–[6] suffer from at least one of five major limitations.

First, detection algorithms invariably assume that the image quality is uniformly excellent across the entire image. However, real data collected at sea often violates this assumption. In practice, images often become noisy at long range where the requirements for SAS processing are more stringent [7]. The degradation of image quality severely impacts the "purity" of the shadows cast by targets, and therefore, also detection performance.

Second, many detection algorithms do not sufficiently exploit the range-dependent nature of target echoes and shadows that can be predicted based on physics-based propagation models and geometrical considerations, respectively. Specifically, it is known *a priori* that the signal-to-noise ratio (SNR) will be weaker at short range, where the grazing angle with the seabed is closer to normal, and that the length of shadows cast by an object will be longer at longer ranges.

Third, most detection algorithms are not tailored to the environmental conditions where the surveying occurs, relying

instead on training data from a different site. By not taking the seabed composition into account, the detection results in certain environments become biased.

Fourth, the detection threshold of a detector is, in general, an arbitrary score with no physical meaning. As a result, there is no principled way to set an appropriate threshold at which to conduct detection operations.

Fifth, most detection algorithms rely on matched-filtering-type approaches that employ discrete Fourier transforms. Although *relatively* fast, these methods require the full sonar image to be available, thereby precluding the possibility of "streaming" real-time detection.

III. Proposed Adaptive Detection Algorithm

The overall objective is to create a robust algorithm that will reliably detect underwater man-made objects of interest. However, the algorithm must also be fast to permit real-time streaming detection onboard an AUV equipped with limited processing capabilities. Since no human intervention is allowed, the algorithm must be flexible enough to sense and adapt to changing environmental conditions from the data collected *in situ*. At the same time, the proposed method should directly address and overcome the limitations that plague existing detection algorithms.

To satisfy these requirements, extensive domain-specific knowledge we possess about the problem is exploited. Emphasis is also placed on tailoring the algorithm, where possible, to the fundamental underlying physics and geometry of the application. The result is a hybrid of rigor and flexibility. Additionally, the cascaded algorithm architecture is designed to minimize computational costs by operating on progressively smaller portions of the image at each stage.

The following sections will describe each step of the detection algorithm that resulted from these considerations. The algorithm description will be supplemented by figures illustrating the effects of each step on an example sonar image, shown in Fig. 1(a), which contains four man-made objects of interest.

A. Integral Image

An integral image [8] is an image representation that allows for very fast computation of rectangular, Haar-like features at any scale or location in constant time. In subsequent stages of the algorithm, such rectangular features will be invaluable for assessing certain distinguishing characteristics of objects, such as echo and shadow levels. The use of an integral-image representation is also a key that makes *streaming* detection possible.

Thus, rather than operating on the pixel-based system of the sonar image, we immediately transform to an integral-image system. In fact, one need not ever form the full sonar image onboard the AUV since the integral-image representation contains equivalent information.

Starting from an original sonar image, A, the corresponding integral image, I, is constructed as follows. The value at a location (x, y) in the integral image corresponds to the sum of the pixels above and to the left of (x, y), inclusive, in the original image, A. That is,

$$I(x,y) = \sum_{x' \leq x, y' \leq y} A(x', y'). \tag{1}$$

The integral image is quickly generated using the recursive relation

$$I(x,y) = I(x-1,y) + z(x,y), \tag{2}$$

where $z(x, y)$ is the cumulative sum of pixels in a row of the original image,

$$z(x,y) = z(x,y-1) + A(x,y). \tag{3}$$

The integral-image approach is particularly well-suited for tasks with streaming sonar data because the construction of an integral image exactly mimics the manner in which the data is collected, namely in a row-wise fashion. For a given along-track position of the AUV, the sonar returns corresponding to a set of ranges — which comprise the columns of the sonar "row-image" — are used in the calculation.

The integral-image representation corresponding to the sonar image in Fig. 1(a) is shown in Fig. 1(b).

Once the integral image is constructed, the ripple detection algorithm never again requires use of the original sonar image. Subsequent stages of the detection algorithm instead repeatedly exploit the integral image.

B. Background Estimation

The first use of the integral image, I, is in the estimation of the sonar-image background map, B. The purpose of the background map is to establish the reverberation level of the seabed in order to subsequently determine locations of shadows in the image.

The reverberation level is a strong function of the seabed composition, so assuming a pre-defined threshold to determine what constitutes background levels is not reliable. For instance, the reverberation level of a soft muddy seabed will be lower than that of a seabed of hard-packed sand. In fact, natural seabed variations can cause the reverberation level to vary substantially at a given site or even within the same image.

A data-driven approach that defines the background level according to some global average over the sonar image could therefore fail catastrophically. Such a method would also require possession of a full image, which would preclude the possibility of streaming, real-time execution. Instead, we argue that the background estimation should be performed in the image *locally*.

Thus, in this work, a robust estimation of the background is obtained by using the local characteristics of the seabed in the measured data. The inherent flexibility afforded by tailoring the estimation to the data itself improves the method's robustness for successfully dealing with different environments. The adaptive estimation also eliminates the possibility of training data mismatch [9], [10], since no training data is used.

The background estimation is calculated using a split window [11] that we call an "Austrian flag" template since it is

composed of three equally-sized horizontal bands of which the top and bottom bands (red in the flag) contribute to the background value, while the middle band (white in the flag) does not. Each band, or rectangle, is of the same size, $b_x = 2.5$ m in the along-track direction and $b_y = 5.0$ m in the range direction.

The background score at location (x, y), $B(x, y)$, is then the mean pixel value in the two rectangles with the "flag" centered around (x, y). Thanks to the integral-image representation, the calculation of the two rectangles that contribute to the background value at a given location can be computed quickly, with a total of only eight array references from the integral image.

Specifically, the background score at (x, y) is calculated as

$$\begin{aligned} B(x,y) = a_b^{-1} \times \\ \Bigg[I\left(x - \frac{3\delta x}{2}, y - \frac{\delta y}{2}\right) - I\left(x - \frac{3\delta x}{2}, y + \frac{\delta y}{2}\right) \\ -I\left(x - \frac{\delta x}{2}, y - \frac{\delta y}{2}\right) + I\left(x - \frac{\delta x}{2}, y + \frac{\delta y}{2}\right) \\ +I\left(x + \frac{\delta x}{2}, y - \frac{\delta y}{2}\right) - I\left(x + \frac{\delta x}{2}, y + \frac{\delta y}{2}\right) \\ - I\left(x + \frac{3\delta x}{2}, y - \frac{\delta y}{2}\right) + I\left(x + \frac{3\delta x}{2}, y + \frac{\delta y}{2}\right) \Bigg], \end{aligned} \quad (4)$$

where δx and δy are the numbers of pixels that correspond to b_x and b_y, respectively, and a_b is the total number of pixels involved in the sums, so that the result is the mean pixel value from the two rectangular areas.

The region immediately surrounding a location (*i.e.,* the "middle band") is not used in the background calculation because those pixels, if an object were present, would be related to the object.

The resulting background map from using the integral image in Fig. 1(b) is shown in Fig. 1(c).

C. Shadow Estimation

Next, the integral image is used again to construct a shadow map. The purpose of the shadow map is to aid in the determination of the locations of shadows in the image that could have been produced by objects of interest.

Recall that the objective of the detection algorithm is to detect objects of interest on the seafloor. The extensive domain-specific knowledge we possess allows us to estimate the minimum height above the seafloor that such an object would extend.

But the unique geometry of the problem — involving the AUV and sonar and objects on the seafloor — provides considerably more information that can be exploited in the detection algorithm. Any object that rises above the seafloor will necessarily cast a shadow (except in pathological cases that can be ignored here). Moreover, the length of the shadow that is cast will be a function of the relative positions of the object and the sonar.

For an AUV operating at an altitude a above the seafloor, and for an object of height h sitting on the seafloor at a range r away from the AUV, simple geometry dictates that the length of the shadow that will be cast by the object will be

$$s_y = \frac{hr}{a - h}. \quad (5)$$

The value of h is set to the minimum height of an object that we are interested in detecting, based on extensive domain knowledge; a is measured directly onboard the AUV.

To respect the inviolable geometry of the problem, the size of the rectangular template used for the shadow calculation is made range-dependent, effectively refining the search for simple shadows to a search for shadows *of an appropriate length*. This choice instills an additional layer of rigor into the method. The size of the rectangle used in the calculation is $s_x = 0.5$ m in the along-track direction, but s_y from (5) in the range-direction.

The shadow score at (x, y) is taken to be the mean pixel value over the rectangle (centered around (x, y)), again calculated quickly thanks to the integral-image representation,

$$\begin{aligned} S(x,y) = a_s^{-1} \times \\ \left[I\left(x - \frac{\delta x}{2}, y\right) - I\left(x - \frac{\delta x}{2}, y + \delta y\right) \right. \\ \left. -I\left(x + \frac{\delta x}{2}, y\right) + I\left(x + \frac{\delta x}{2}, y + \delta y\right) \right], \end{aligned} \quad (6)$$

where δx and δy are the numbers of pixels that correspond to s_x and s_y, respectively, and a_s is the total number of pixels involved in the sum, so that the result is the mean pixel value from the rectangular area.

The resulting shadow map from using the integral image in Fig. 1(b) is shown in Fig. 1(d).

It should be noted that because of the integral-image representation, the range-dependent nature of the template causes no additional computational complexity. In contrast, existing detection algorithms that employ matched-filtering-type methods that rely on discrete Fourier transforms cannot enjoy this range-dependent flexibility. Instead, those approaches will be inherently mismatched (since the geometry of the problem is not respected) and inferior; those approaches must either use a single average range (which biases against detecting targets at certain ranges) or repeat the whole process multiple times for different expected shadow sizes, which increases the computational time. Also, the approaches employing Fourier transforms require a full image, which precludes streaming detection.

D. Addressing Poor Image Quality

It is commonly assumed that all sonar images are of good quality everywhere, but that is often not the case with real sonar data collected at sea [12]. In particular, image quality often degrades significantly at long range where the effects of multipath manifest [13], or where the more stringent motion correction requirements for successful SAS image formation cannot be satisfied [7].

In this poor image quality regime, shadow purity is diminished (*i.e.,* shadows "fill in"), making shadows an unreliable

clue for object detection. Therefore, in the proposed algorithm, all regions in the image where the image quality is poor are automatically retained to be examined further in the subsequent stages.

The image quality is quantified by the peak correlation between successive pings at a given range (*i.e.,* distance from the sonar) since this quantity, termed the coherence, is directly proportional to the image's SNR [14], and has been shown [15] to be a reliable predictor of SAS image quality.

E. Region of Interest Determination

After a background map, B, and a shadow map, S, are constructed, the regions of the image that may possibly contain targets of interest are determined.

Any pixel for which the shadow map value is sufficiently lower than the corresponding background map value is declared to be a region of interest that will receive further investigation. Specifically, if $S(x,y) < \beta B(x,y)$, then the pixel (x,y) is considered to be part of a shadow and therefore a region of interest that will receive further investigation. The scaling factor β makes the requirement for shadows more stringent. (Essentially, this *defines* that a shadow is when the pixel value is a certain amount less than the average value of the surrounding background.)

This test is the first data-reduction stage of the detection cascade. It achieves a large-scale reduction in pixels that must be examined further — in practice, more than 90% of the pixels are usually removed in this step — thereby greatly reducing computation loads of all subsequent stages.

The binary result of the background and shadow map comparison is shown in Fig. 1(e), where white pixels correspond to regions of potential alarms that will be investigated further.

It should be noted that we do not simply search for low pixel values in the sonar image. Rather, we tailor the detection algorithm to mimic what we hypothesize the human visual system does when performing detection: searching for pixels or groups of pixels for which the *relative* pixel value is lower than the values of the surroundings. That is, we exploit contextual information, via the background map, in the detection process.

In addition, the algorithm is made more robust because it focuses on an intrinsic aspect of the problem: objects of interest must rise a certain height above the seafloor, which means a shadow of a particular length will be cast. Since naturally occurring highlights (not caused by objects extending above the seabed) in sonar images are common, searching first for locations of high signal strength could generate many false alarms.

F. Echo Estimation

Next, the echo score is calculated for the remaining potential alarms by again exploiting the integral image. Specifically, the mean pixel value in a rectangular area about a given location is used as the echo score.

The size of the rectangle used in the calculation is $e_x = 0.75$ m in the along-track direction and $e_y = 0.5$ m in the range direction. This size is chosen because it is related to the smallest object of interest we wish to detect.

The echo score at (x,y) is taken to be the mean pixel value over the rectangle, calculated as

$$E(x,y) = a_e^{-1} \times \left[I\left(x - \frac{\delta x}{2}, y - \frac{\delta y}{2}\right) - I\left(x - \frac{\delta x}{2}, y + \frac{\delta y}{2}\right) - I\left(x + \frac{\delta x}{2}, y - \frac{\delta y}{2}\right) + I\left(x + \frac{\delta x}{2}, y + \frac{\delta y}{2}\right) \right], \quad (7)$$

where δx and δy are the numbers of pixels that correspond to e_x and e_y, respectively, and a_e is the total number of pixels involved in the sum, so that the result is the mean pixel value from the rectangular area.

It should be noted that a minor realignment must be effected prior to the echo score calculations because the potential alarm locations are located in the shadows that are cast by the object that is nearer in range (*i.e.,* closer to the sensor), while the echo scores that we desire should correspond to the highlights produced by the objects themselves (not their shadows). Therefore, the locations at which the echo score is computed are all shifted closer in range by $3e_y/2$.

G. Echo-Loss Correction

The incident angle of the sonar signal is closer to normal at shorter ranges, so the seabed will reflect back more energy to the sonar receiver on the AUV. Natural propagation loss means that the sonar returns from the seabed at longer ranges will be weaker. As a result, the SNR of a target at short range will necessarily be lower than the SNR of the same target at longer range. After standard range-normalization is performed to achieve approximately range-independent pixel levels in the sonar image, the echoes at shorter ranges will effectively be decreased.

To eliminate this undesirable natural phenomenon, which would make targets at shorter range more difficult to detect, a rigorous range-dependent correction term that is based on the well-known propagation loss equation [16] is applied. The echo correction that we add to the echo score of a potential alarm at range r is

$$\epsilon(r) = 20 \log_{10}\left(\frac{r_0}{r}\right), \quad (8)$$

where r_0 is a reference range set to the nominal maximum range of the sonar (here, $r_0 = 150$ m).

The resulting echo map is shown in Fig. 1(f), where it can be observed that the calculations are performed only at the potential alarm locations, which again saves computation.

H. Region of Interest Determination Redux

The echo correction term ensures that the echo score of a given target will be approximately range independent. In turn, this permits the use of a single detection threshold for all ranges.

The echo scores are directly related to the signal strength of the objects, which means that setting a detection threshold

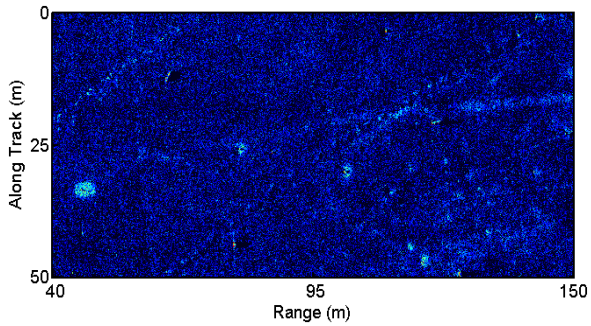

(a) SAS image

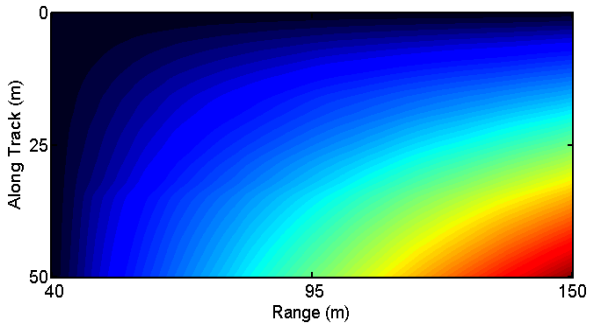

(b) Integral image

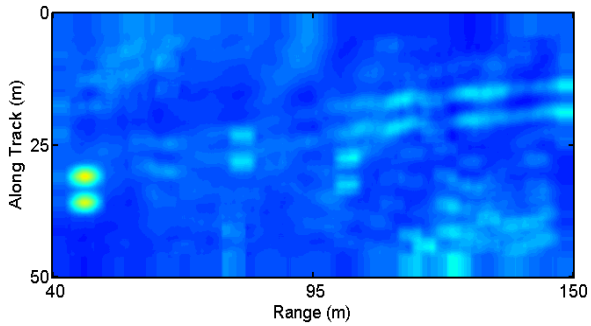

(c) Background map

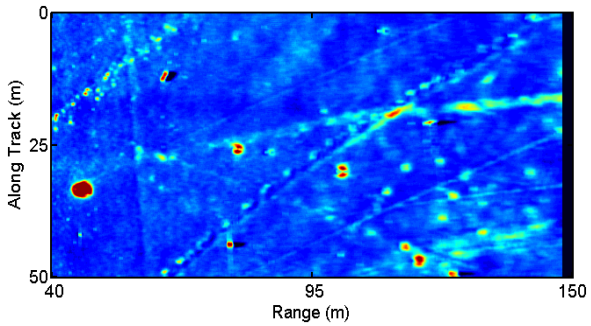

(d) Shadow map

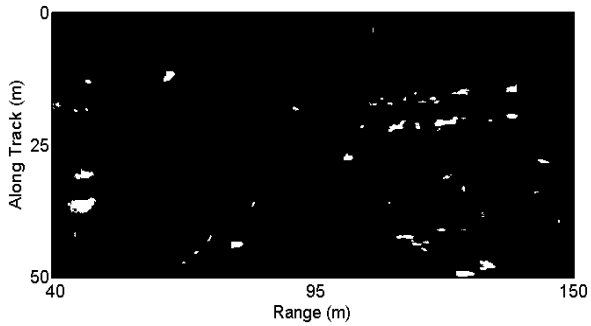

(e) Potential alarms after shadow assessment

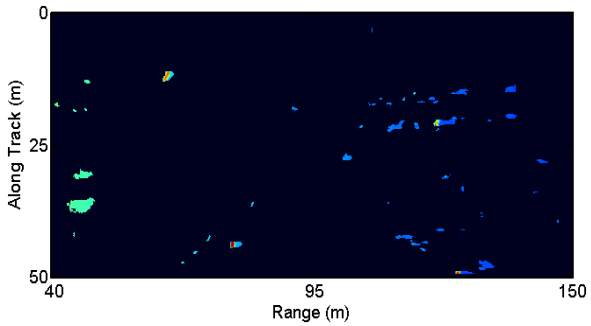

(f) Echo map

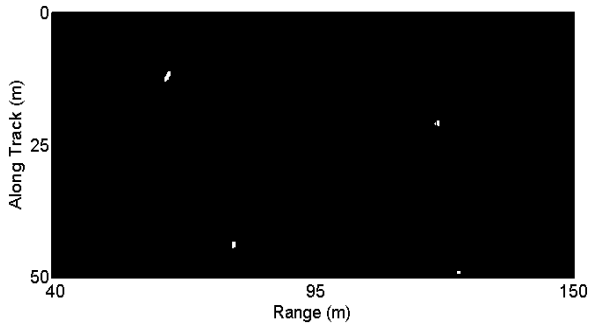

(g) Potential alarms after echo assessment

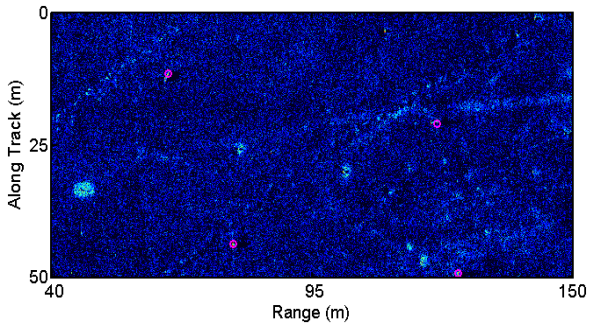

(h) Final detection map

Fig. 1. Detection algorithm. The original sonar image (a) is converted into its equivalent integral-image representation (b), which is then used to estimate a background map (c) and a shadow map (d), the comparison of which generates a map of potential alarms (e). The integral image is again used to estimate the echo map (f) at locations of potential alarms, after which a threshold is applied to obtain the final areas of potential alarms (g), which are then converted into discrete detections (h).

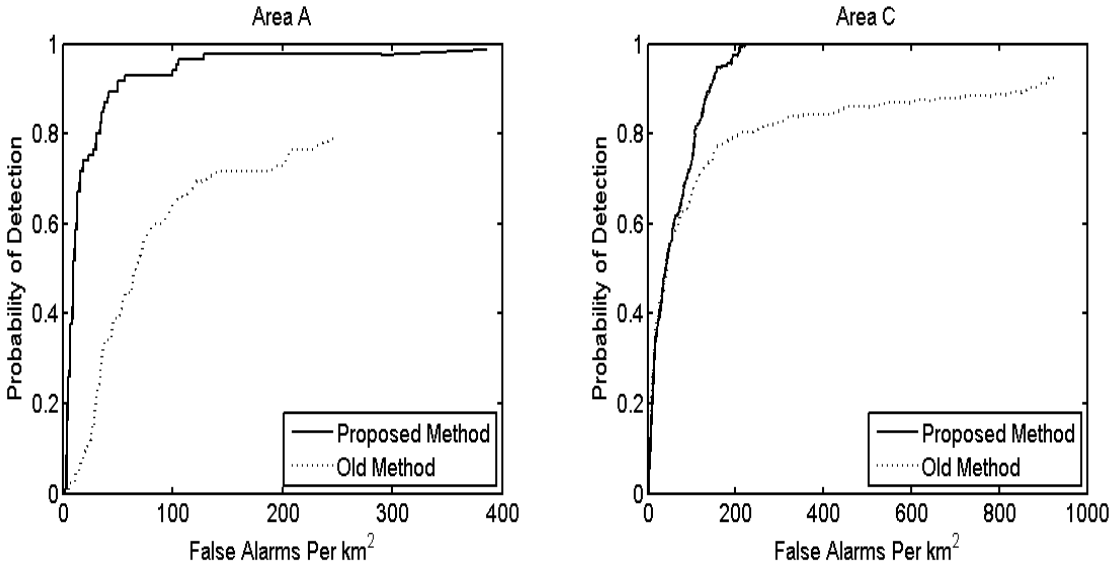

(a) Area A: soft mud

(b) Area C: hard-packed sand

Fig. 2. ROC curves for the proposed detection algorithm compared to the standard (old) detection method at two different locations.

can be determined rigorously. Specifically, the threshold can be set such that we wish to detect any object for which the signal strength exceeds a given level.

In contrast, other detection algorithms are forced to set an arbitrary threshold value because the detection score, usually related to how well an area of the image matches a highlight-shadow pattern characteristic of objects, cannot be equated to a well-defined, tangible, physical quantity.

The final stage of the detection cascade removes those areas for which the echo score is below the desired signal-strength threshold. The map of potential alarms that remain after this stage are shown in Fig. 1(g).

I. Final Detection Map

The final step of the detection algorithm is to convert the areas of potential alarms to a list of discrete alarms. To do this, the location of the maximum echo score — which subsequently is treated as that alarm's detection score — within each potential alarm region is treated as the alarm location.

To present this result visually for our example image, we plot the alarm locations overlaid on the original sonar image, in Fig. 1(h). It can be seen that all 4 targets were successfully detected while incurring zero false alarms. For purposes of comparison, application of a widely used method [6] that employs a matched-filtering-type procedure on this image resulted in 3 false alarms and also failed to detect 2 objects of interest.

IV. Experimental Results

In April-May 2008, the NATO Undersea Research Centre (NURC) conducted the Colossus II sea trial in the Baltic Sea off the coast of Latvia. During this trial, high-resolution sonar data was collected by the MUSCLE AUV, which is equipped with a 300 kHz sonar with a 60 kHz bandwidth that can achieve image resolution of approximately 3 cm.

At two sites with different seabed characteristics, one in Rīga Bay ("Area A") and one off the coast of Liepāja ("Area C"), a set of man-made objects of interest (of different shapes and characteristics) were deployed and a series of AUV surveys was performed over the area. A summary of the data set details is shown in Table I. The seabed at Area A is softer, with a muddy composition. The seabed at Area C is characterized by hard-packed sand.

TABLE I
Data set details

Data Set	Seabed Composition	Object Looks	Area (sq. km)
Area A	Soft Mud	85	1.7474
Area C	Hard Sand	224	2.3207

To demonstrate the promise of the proposed detection algorithm, we evaluated its performance on these data sets. For purposes of comparison, we also considered the standard detection algorithm used in [6], since slight variations of this popular method are in widespread use. (This comparison is also fair because neither of the methods exploit or require any training data.) This standard ("old") method, which suffers from the various limitations noted previously, performs detection by correlating a template, consisting of a generic highlight-shadow pattern characteristic of man-made objects of interest, with the sonar images.

The resulting detection performance at the two sites is shown in Fig. 2 in terms of receiver operating characteristic (ROC) curves, generated by varying the final detection thresh-

old. It can be seen that the proposed method achieves superior performance to the popular existing method.

The gain in performance at Area A, which was characterized by soft mud, can be attributed to the proposed algorithm's ability to detect objects with weaker responses. In particular, the echo correction term ensures that objects at short range will be detected at the same rate (and at comparable detection scores) as objects at long range.

The gain in performance at Area C, which was characterized by hard-packed sand, is because the proposed algorithm bases its initial search on shadows rather than strong echoes. Naturally occurring variation on the seabed that can generate strong echoes is common in hard-packed sand; such areas, even if lacking a strong shadow, then necessarily generate high detection scores using the old method, since the contribution of large-valued highlight pixels overwhelms that of shadow pixels in the template score.

It should also be noted that the proposed algorithm's use of an integral-image representation and a cascaded architecture allow very quick execution such that streaming real-time detection onboard an AUV would be feasible.

V. Discussion

The five previously mentioned limitations that plague many existing detection algorithms were taken into account during the construction of the proposed algorithm. As a result, the new method overcomes all of these limitations.

First, the proposed detection algorithm explicitly takes the image quality into account by relaxing the requirements for the shadow purity where the image quality, quantified by the peak correlation between successive pings, is poor.

Second, the algorithm defines the size of the shadow template based on geometrical considerations, adapting based on the altitude of the AUV, so that the template size is a function of range. (Because of the new manner in which the algorithm is constructed, the detection can still be performed over an entire image in a single computation, even with range-dependent template sizes.) In addition, the detection scores are also weighted according to a well-known, physical propagation-loss model, so that near-range objects are not systematically biased with lower detection scores. Instead, a given object will produce detection scores that are independent of range.

Third, the algorithm is intimately tied to the environmental conditions by taking into account the composition of the seabed via the local estimation of the background. As a result, shadows can still be reliably detected even in soft seabed, which will have a weaker reverberation level that is closer in value to an actual shadow.

Fourth, the algorithm is formulated such that the detection score is directly proportional to the signal strength of an object. Therefore, setting the detection threshold has real, physical meaning: the minimum signal strength of an object for which detection is mandated.

Fifth, the proposed algorithm employs a remarkable construct from the vision community, the integral-image representation of the sonar image, that makes fast, real-time detection possible by performing a matched-filtering-like procedure using only a handful of addition operations.

VI. Conclusion

A fast, novel algorithm for the detection of underwater man-made objects was presented. The promising proposed approach addresses and overcomes several limitations of existing underwater object detection algorithms. By exploiting integral-image representations and by adhering to a cascaded architecture, the algorithm makes streaming, near-real-time detection onboard an AUV feasible.

Related ongoing work that will be continued in the future is concentrated on reducing false alarms caused by sand ripples. Preliminary results of a ripple-detection algorithm that can easily be incorporated into the overall object-detection algorithm for this purpose are promising.

Future work will attempt to perform near-real-time detection and survey route adaptation, based on detected objects and sand ripples, in a completely autonomous manner at sea with an AUV.

References

[1] M. Hayes and P. Gough, "Broad-band synthetic aperture sonar," *IEEE J. Oceanic Engineering*, vol. 17, no. 1, pp. 80–94, 1992.

[2] G. Dobeck, J. Hyland, and L. Smedley, "Automated detection/classification of seamines in sonar imagery," in *Proc. SPIE International Society of Optics*, vol. 3079, 1997, pp. 90–110.

[3] S. Reed, Y. Petillot, and J. Bell, "An automatic approach to the detection and extraction of mine features in sidescan sonar," *IEEE Journal of Oceanic Engineering*, vol. 28, no. 1, pp. 90–105, January 2003.

[4] J. Fawcett, A. Crawford, D. Hopkin, V. Myers, and B. Zerr, "Computer-aided detection of targets from the CITADEL trial Klein sonar data," Defence R & D Canada - Atlantic, Canada, Tech. Rep. DRDC Atlantic TM 2006-115, 2006.

[5] F. Maussang, J. Chanussot, A. Hétet, and M. Amate, "Mean-standard deviation representation of sonar images for echo detection: Application to SAS images," *IEEE J. Oceanic Engineering*, vol. 32, no. 4, pp. 956–970, 2007.

[6] J. Groen, E. Coiras, and D. Williams, "Detection rate statistics in synthetic aperture sonar images," in *Proc. Intl. Conf. & Exh. Underwater Acoustic Measurements*, 2009, pp. 367–374.

[7] A. Bellettini and M. Pinto, "Theoretical accuracy of synthetic aperture sonar micronavigation using a displaced phase-center antenna," *IEEE J. Oceanic Engineering*, vol. 27, no. 4, pp. 780–789, 2002.

[8] P. Viola and M. Jones, "Robust real-time object detection," *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2004.

[9] G. Widmer and M. Kubat, "Learning in the presence of concept drift and hidden contexts," *Machine Learning*, vol. 23, pp. 69–101, 1996.

[10] B. Zadrozny, "Learning and evaluating classifiers under sample selection bias," in *Proceedings of the 21st International Conference on Machine Learning*, 2004.

[11] H. Cox and D. Pace, "A fast normalizer," in *Thirtieth Asilomar Conference on Signals, Systems and Computers*, vol. 1, 1996, pp. 459–463.

[12] R. Hansen, H. Callow, T. Sæbø, and S. Synnes, "Challenges in seafloor imaging and mapping with synthetic aperture sonar," in *Proceedings of European conference on Synthetic Aperture Radar*, 2010.

[13] L. Wang, G. Davies, A. Bellettini, and M. Pinto, "Multipath effect on DPCA micronavigation of a synthetic aperture sonar," in *Impact of Littoral Environmental Variability on Acoustic Predictions and Sonar Performance*, N. Pace and F. Jensen, Eds. Kluwer Academic Publishers, 2002, pp. 465–472.

[14] S. Synnes, R. Hansen, and T. Sæbø, "Assessment of shallow water performance using interferometric sonar coherence," in *Proceedings of International conference on Underwater Acoustic Measurements*, 2009.

[15] D. Williams, "Image-quality prediction of synthetic aperture sonar imagery," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2010, pp. 2114–2117.

[16] R. Urick, *Principles of Underwater Sound*. McGraw-Hill, 1983.