

On Adaptive Underwater Object Detection

David P. Williams

Abstract—A new algorithm for the detection of underwater objects in sonar imagery is proposed. One particularly novel component of the algorithm also detects the presence of, and estimates the orientation of, sand ripples. The overall 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 in streaming sonar data collected by an autonomous underwater vehicle (AUV) 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. Moreover, the proposed algorithm achieves superior performance across a variety of seabed types on a large, challenging data set of real sonar data collected at sea. Ways to exploit the findings and adapt AUV surveys for optimized detection performance are also suggested.

I. INTRODUCTION

A. Motivation

The high-resolution imaging of underwater environments afforded by sonar has proven particularly useful for the detection of 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 such missions, 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.

To achieve this goal, two major obstacles must be overcome. First, 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. Second, a plan is needed for specifying how the information gleaned from the detection results can be exploited to intelligently adapt the AUV route.

This work addresses both challenges, and in the process, removes two of the most significant barriers preventing fully autonomous detection missions using an AUV.

B. Contribution and Relevance of Paper

The principal contribution of this paper is a new algorithm for the detection of underwater objects in sonar imagery. The

flexible yet rigorous approach overcomes several limitations of existing detection algorithms that are in widespread use. Additionally, the algorithm’s cascaded architecture ensures that it can be executed very quickly, making real-time detection in streaming sonar imagery collected onboard an AUV during a survey feasible.

An auxiliary contribution of the work is a novel application of the integral-image formulation for the detection of sand ripples and the estimation of their orientation. This ripple detection algorithm also provides insight into how AUV routes can be adapted *in situ* to optimize data collection and mitigate the concealment effects of sand ripples on object detection. As such, this paper advances the field toward intelligent, adaptive, completely autonomous underwater surveys without any “human in the loop.”

The ideas in this work can also be exploited to ameliorate the problem of “data deluge” and reduce operator fatigue [1], and to dramatically speed up mission times (*e.g.*, by eliminating the need to redeploy the AUV to collect additional or better data).

Although the primary focus of the work here is for the task of generic object detection, the algorithms can also be used for a wide range of other marine robotics applications, including habitat mapping [2], seabed classification [3], archaeology [4], and pipeline monitoring [5]. The methods proposed are also relevant to other fields that employ airborne sensors for image-processing-based tasks. For example, the fast detection framework can be exploited for ground-based target-detection [6], while the ripple detection algorithm would be useful for oil exploration applications that require the detection of sand dunes [7].

C. Organization of Paper

The remainder of this paper is organized as follows. Sec. II notes the limitations of existing underwater object detection algorithms and outlines an overview of the proposed approach. The three stages of the proposed cascaded detection algorithm — shadow detection, ripple detection, and echo detection — are described in Sec. III, Sec. IV, and Sec. V, respectively. Experimental results on real, measured sonar data collected at sea are presented in Sec. VI. A discussion outlining how the detection results can be exploited for the *in situ* adaptation of AUV survey routes is provided in Sec. VII. Concluding remarks are made in Sec. VIII.

II. UNDERWATER OBJECT DETECTION

A. Limitations of Existing Detection Algorithms

Currently, after sonar data has been collected, an object detection algorithm is applied post-mission. Typically, these

D. Williams is with the NATO Undersea Research Centre, La Spezia, Italy. williams@nurc.nato.int

detection algorithms search for highlight-shadow patterns characteristic of objects [8]–[12]. 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.

However, existing object detection algorithms suffer from five major limitations. First, detection algorithms invariably assume, incorrectly, that the image quality is uniformly excellent across the entire image. Second, many detection algorithms do not sufficiently exploit the range-dependent nature of object echoes and shadows that can be predicted based on physics-based propagation models and geometrical considerations, respectively. Third, most detection algorithms are not tailored to the environmental conditions where the surveying occurs, relying instead on training data from a different site. Fourth, the detection threshold of a detector is, in general, an arbitrary score with no physical meaning. Fifth, most detection algorithms rely on matched-filtering-type approaches that employ discrete Fourier transforms and preclude the possibility of “streaming” real-time detection.

B. Overview of Proposed Algorithm

The proposed object detection algorithm is composed of three major stages: shadow detection, ripple detection, and echo detection. We first briefly provide an overview of the motivations and general philosophy used to create the algorithm.

The overall objective is to create a robust algorithm that will reliably detect underwater 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 — briefly, due to space constraints — each step of the detection algorithm that resulted from these considerations. But the algorithm description will be supplemented by figures illustrating the effects of each step on one particularly challenging sonar image, shown in Fig. 1(a). (For reference, the left portion of the image is characterized by sand ripples, while the right portion is flat sandy seabed. Seven objects of interest are present in the image.)

III. SHADOW DETECTION

The first stage of the proposed object detection algorithm establishes locations of shadows in the image. To do so, the sonar image is first converted to an integral-image representation. This image is then used to quickly estimate a background map and a shadow map, the comparison of which provides locations of shadows (consistent with objects of interest) that are examined further in the second stage of the algorithm.

A. Integral Image

An integral image [13] is an image representation that allows for very fast computation of rectangular, Haar-like features at any scale or location in constant time. 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* object detection possible. In fact, 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.

Thus, rather than operating on the pixel-based system of the sonar image, we immediately transform to an integral-image system.¹ The integral-image representation corresponding to the sonar image in Fig. 1(a) is shown in Fig. 1(b).

B. Background Estimation

The reverberation level of the seabed is a strong function of its composition (*e.g.*, soft mud or hard sand), so assuming a pre-defined threshold to determine what constitutes background levels is not reliable. In fact, natural seabed variations can cause the reverberation level to vary substantially at a given site or even within the same image. Therefore, we estimate the background in the image using the *local* characteristics of the seabed.

Specifically, the background score associated with a pixel is taken to be the mean pixel value of a fixed-size rectangular template centered around the point, computed quickly by exploiting the integral image. The resulting background map from using the integral image in Fig. 1(b) is shown in Fig. 1(c).

The flexibility afforded by tailoring the estimation to the measured *in situ* data makes the method robust and successful in different environments. This adaptive estimation also eliminates the possibility of training data mismatch [14], since no training data is used.

C. Shadow Estimation

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.

¹In fact, one need not ever form the full sonar image onboard the AUV since the integral-image representation contains equivalent information.

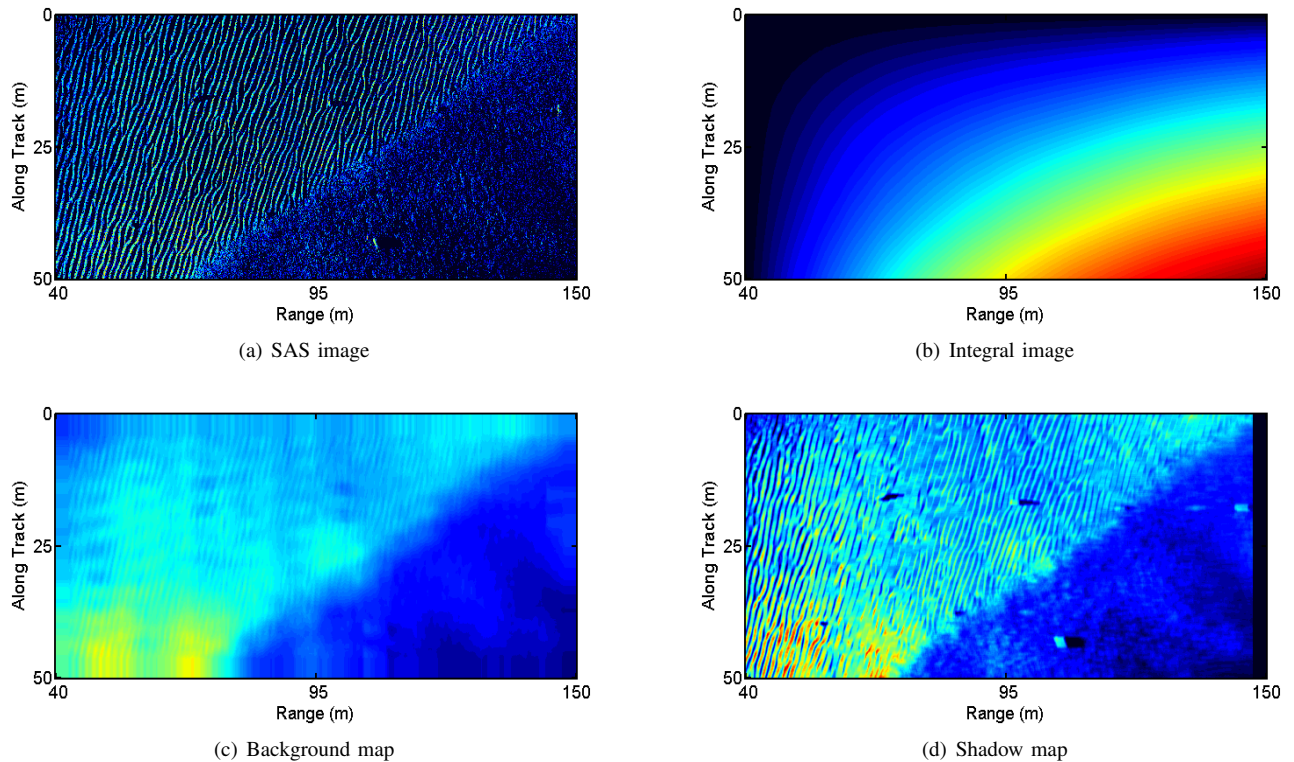


Fig. 1. Detection algorithm, part 1 of 3. The original sonar image (a) is converted into its equivalent integral-image representation (b), which is then used to estimate a background map (c) and a shadow map (d).

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 $s = hr/(a - h)$. The value of h is set to the minimum height of an object of interest, based on extensive domain knowledge; a is measured directly onboard the AUV.

Therefore, the shadow estimation is refined to instead be a search for shadows of an appropriate length (*i.e.*, that could have been produced by objects of interest) by exploiting this geometry.

The shadow score at a pixel is taken to be the mean pixel value over a *range-dependent* rectangle centered around the point, again calculated quickly thanks to the integral-image representation. 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 and rely on discrete Fourier transforms cannot enjoy this range-dependent flexibility (and therefore ignore the inviolable geometry of the problem).

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 [15]. In particular, image quality often degrades significantly at long range where the effects of

multipath manifest [16], or where the more stringent motion correction requirements for successful SAS image formation cannot be satisfied [17].

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 the peak correlation of consecutive ping returns, termed the coherence, as a function of range, since this measure is directly related to the signal-to-noise ratio (SNR) of an image [18].

E. Region of Interest Determination (Stage I)

Any pixel for which the shadow map value is sufficiently lower than the corresponding background map value is declared to be a shadow that will receive further investigation.

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. 2(a), where white pixels correspond to regions of potential alarms that will be investigated further. The vertical stripe of white at longer range is due to poor image quality at those ranges.

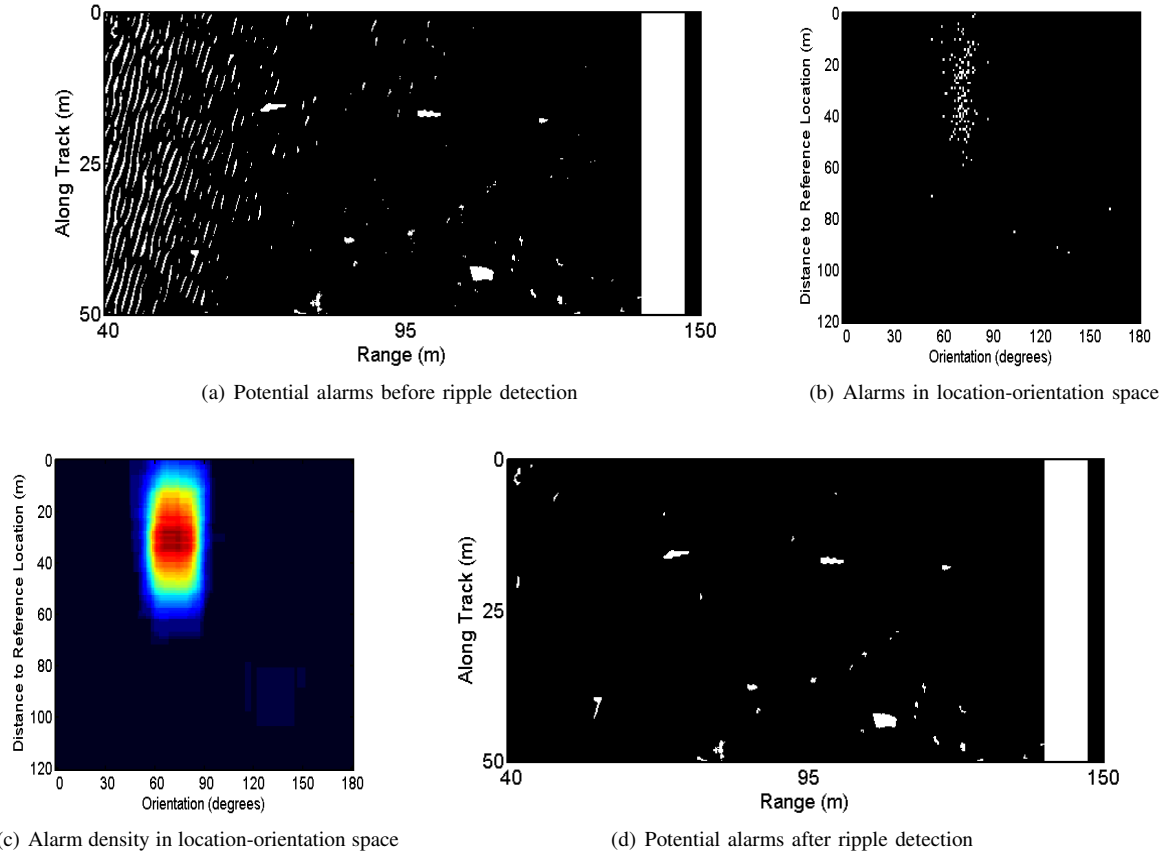


Fig. 2. Detection algorithm, part 2 of 3. The result of the background map and shadow map comparison produces a map of potential alarms (a). A location and orientation feature are then extracted for each alarm (b), from which density estimation is performed (c) to determine the presence and orientation of sand ripples. Alarms deemed to be sand ripples are removed from further consideration (d).

IV. RIPPLE DETECTION

The second stage of the proposed object detection algorithm detects ripples in the sonar image. A feature extraction step maps each shadow to a point in a new space in which ripples, characterized by a high density of elongated shadows that are oriented similarly, can be efficiently detected. Locations that do not belong to ripples are examined further in the third stage of the algorithm.

A. Overview

Sand ripples are one of the most confounding factors and frequent sources of false alarms for underwater object detection because they produce highlight-shadow patterns — the sizes of which will depend on the wavelength and range of the ripples — characteristic of some objects. Additionally, when objects lie in ripple fields, the echoes of the objects and ripples blend together, as do their respective shadows, making detection challenging.

Since the overall objective here is to detect objects, the ripple detection algorithm we propose is more accurately called a *nuisance* ripple detection algorithm. Rather than attempting to detect all sand ripples in an image, the algorithm instead attempts to detect only those ripples that affect our ability to

detect objects, *i.e.*, nuisance ripples.²

B. Motivation

To overcome the challenges posed by sand ripples, we develop a principled ripple detection algorithm that is tailored both to the fundamental characteristics that *define* sand ripples and the underlying physics that creates them.

Sand ripples are formed when currents or waves flow in a particular direction [19]. Because of this mechanism by which they are created, ripples tend to exist in large *fields*, rather than as individual instantiations.

Although ripples will be characterized by a dominant orientation (dictated by the direction of flow), considerable variability will also exist. Ripples will in fact span a range of orientations, with varying amplitudes and periods as well.

Additionally, the relative geometry between the sonar and the mound of sand that defines the ripple means that the shadows that are cast by the ripples will be elongated along the direction of the ripple crest (or trough) [20].

These three fundamental physical characteristics of ripples, which can be observed in Fig. 1(a), will be exploited in the ensuing stages of the ripple-detection algorithm.

²In order to instead detect *all* ripples, such as for environmental assessment purposes, the template used in the shadow map determination should be made range-independent.

C. Feature Extraction

To efficiently search for a high density of elongated shadows that are oriented similarly, each discrete shadow in the binary shadow map (or potential-alarm map, cf. Fig. 2(a)) is mapped into a new feature space in \mathbb{R}^2 in which it is easy to detect such characteristics.

Specifically, the centroid and orientation are calculated for each discrete shadow (*i.e.*, each connected region, or “blob”) in the potential-alarm map. The orientation of a shadow is obtained by fitting an ellipse to the shadow area. The centroid is converted to a location feature that is simply the distance of the centroid to a fixed reference location in the mission (or image).

The result of this feature extraction process that transforms the shadows in sonar-image space in Fig. 2(a) to location-orientation space is shown in Fig. 2(b).

D. Density Estimation

By discretizing this new feature space and treating it as a two-dimensional “image,” the corresponding integral image can be computed readily. The integral-image formulation can then be exploited to quickly perform density estimation in this space.

Specifically, the mean number of discrete shadows in a rectangular area about a given location is used as the shadow density estimate. The size of the rectangle used in the calculation is chosen to respect the physical proximity and orientation variability of the underlying mechanisms that help define ripple fields.

The resulting density estimate is shown in Fig. 2(c).

If the maximum density is above a predefined threshold, sand ripples are declared to be present in the image. This threshold effectively *defines* our concept of a ripple field, by establishing the minimum density of potential alarms at similar orientations needed to declare a ripple field.

The principal (center) orientation of the ripple field is estimated to be the orientation at which the density achieves a maximum. The ripple field is deemed to span a set of orientations, with the minimum and maximum orientations of that span set to the orientations at which the density drops to half the maximum density. That is, the span is adaptively estimated automatically based on the evidence that is provided by the data itself.

For the image related to Fig. 2(c), the principal ripple orientation was deemed to be $\theta_r = 73^\circ$, spanning the orientations $\Theta_r = [57^\circ, 86^\circ]$.

E. Region of Interest Determination (Stage II)

If sand ripples have been detected, those potential alarms for which the orientation is within the estimated span of ripple orientations, Θ_r , are removed from further consideration.

The result of the nuisance ripple detection stage is shown in Fig. 2(d), from which it can be observed that the ripples have largely been removed from further consideration.

This test is the second data-reduction stage of the detection cascade, and further reduces the number of pixels that must be examined further.

V. ECHO DETECTION

The third stage of the proposed object detection algorithm estimates the signal strength, or echo level, of the remaining alarms. A range-dependent correction term to counteract natural propagation loss is also applied before determining the final map of object alarms.

A. Echo Estimation

The echo score is calculated for the remaining potential alarms by once again exploiting the integral image. Specifically, the mean pixel value in a rectangular area about a given location³ is used as the echo score.

The size of the rectangle is related to the smallest object of interest. By using a generic rectangular template, a wide range of objects can be detected. In contrast, a template specially tailored to a particular object would not generalize well, and therefore, would fail to detect new but related objects.

B. 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 an object at short range will necessarily be lower than the SNR of the same object 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 objects at shorter range more difficult to detect, a rigorous range-dependent correction term that is based on the well-known propagation loss equation [21] is applied. The echo correction that we add to the echo score of a potential alarm at range r is $\epsilon(r) = 20 \log_{10}(r_0/r)$, 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. 3(a), where it can be observed that the calculations are performed only at the potential alarm locations, which again saves computation.

C. Region of Interest Determination (Stage III)

The echo correction term ensures that the echo score of a given object 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 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

³A shift is applied so the echo calculation is performed over the area preceding the shadow location, since an object that casts a shadow will necessarily be located closer to the sonar.

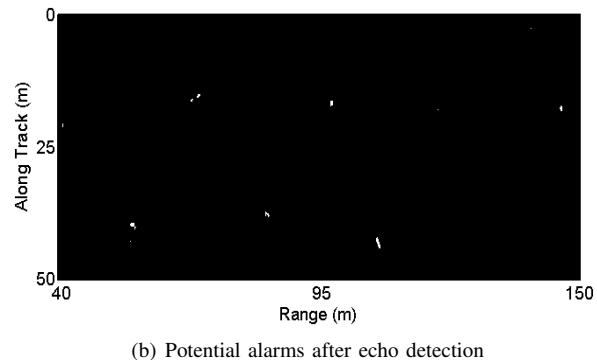
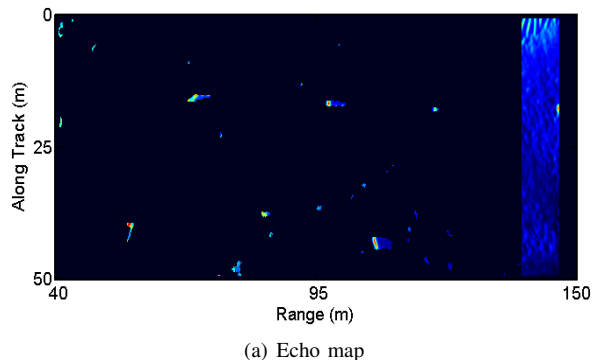


Fig. 3. Detection algorithm, part 3 of 3. A threshold is applied to the echo map (a) to generate the final areas of potential alarms (b), which are subsequently converted into discrete detections.




Fig. 4. Final detection map. White circles indicate false alarms, while red circles mark correct object detections. The size of each circle is proportional to the alarm's detection score.

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 third and final stage of the detection cascade removes those areas for which the echo score is below the desired signal strength threshold. The potential alarms that remain after this stage are shown in Fig. 3(b).

D. 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. 4. It can be seen that all 7 objects were successfully detected (and 6 of the 7 objects were detected while incurring zero false alarms) and 5 false alarms were recorded.

For purposes of comparison, application of a widely used method [12] that employs a matched-filtering-type procedure, on this image resulted in 112 false alarms — most generated by the sand ripples — and also failed to detect 3 objects (two in the ripple field and one at long range where the image quality is poor).

VI. EXPERIMENTAL RESULTS

A. Data Sets

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 each of three sites with different seabed characteristics, one in Rīga Bay (“Area A”) and two off the coast of Liepāja (“Area B” and “Area C”), a set of objects (of different shapes and characteristics) were deployed and a series of AUV surveys was performed over the area. Across the three sites, the average number of object looks was 150 and the average area of seabed surveyed was 2.24 km².

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 [12], 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 objects of interest, with the sonar images.

B. General Results

The resulting detection performance in the three areas is shown in Fig. 5 in terms of receiver operating characteristic (ROC) curves, generated by varying the final detection threshold. 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.

At Area B, the worse false alarm rate of the the old method can be attributed to its inability to reject ripples; the worse

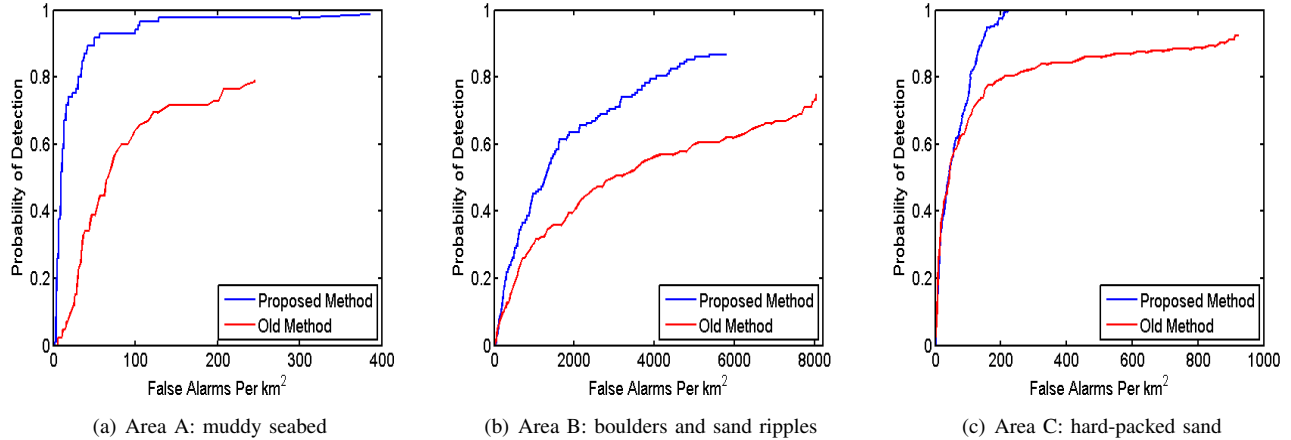


Fig. 5. Detection performance for the three areas in terms of ROC curves. The seabed characteristics of each data set are summarized in the captions.

detection rate is due to its difficulty of detecting objects where the image quality is poor or in sand ripples (since the template correlation-score is skewed by background ripples).

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.

To highlight the appealing data-reduction aspect of the cascaded detection approach, Table I presents the average fraction of an image that is passed on to the echo determination stage of the proposed algorithm. As can be seen from the table, significant computational savings are obtained by employing the cascaded approach.

TABLE I
FRACTION OF IMAGE PASSED ON TO ECHO DETERMINATION STAGE
(MEAN \pm ONE STANDARD DEVIATION, ACROSS DATA SET).

| DATA SET | $\mu \pm \sigma$ |
|----------|---------------------|
| AREA A | 0.0416 ± 0.0591 |
| AREA B | 0.1580 ± 0.1159 |
| AREA C | 0.0074 ± 0.0127 |

It should also be reiterated that the proposed algorithm can be executed very quickly such that streaming real-time detection onboard an AUV would be feasible.

C. Ripple-Orientation Dependence

Unlike in Areas A and C, the proposed algorithm was unable to achieve (nearly) perfect detection in Area B. To investigate the reason for this, we examine the detection performance (of the objects located in the ripple field) as a function of the AUV survey direction. (During the mission, the AUV surveyed in a “lawnmower” fashion in four principal directions; the sonar look-direction (*i.e.*, the

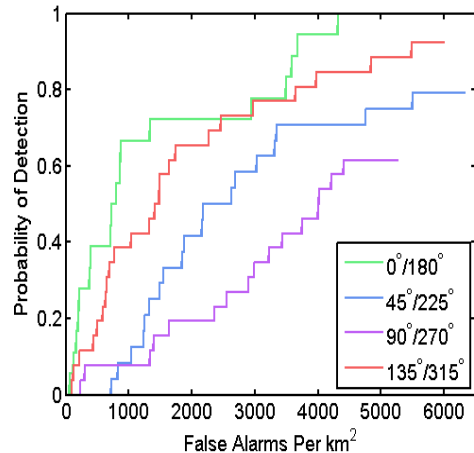


Fig. 6. Detection performance for objects in sand ripples in Area B, as a function of AUV survey direction (legend indicates sonar look-direction). Sand ripples in the area were oriented at $73^\circ/253^\circ$.

range direction) is orthogonal to the direction in which the AUV travels.)

The ROC curves when the performance is presented in this manner is shown in Fig. 6. As can be seen, there is a strong orientation dependence on the detection performance. For example, the results indicate that for a false alarm rate of 1000 false alarms per km^2 , the probability of detection would be 0.67 if the survey was executed with a 0° look-direction, but it would be only 0.08 if the survey was executed with a 90° look-direction.

It is worth noting that by examining the results of the ripple detection stage, it can be established that the approximate principal ripple angle in this area was 73° (visual verification of the images confirms the accuracy of the estimation; *e.g.*, the image in Fig. 1(a) was from this area).

This insight reveals that the further the sonar look-direction was from the ripple orientation, the better the detection performance was, with this trend holding for all four survey directions in Fig. 6. (Therefore, a survey at

163°/343° would have been optimal.)

We hypothesize that this phenomenon arises because when the ripple orientation is similar to the look-direction of the sonar, the shadows associated with ripples appear nearly horizontal, similar to the shadows cast by objects of interest. When an object is present, its shadow more naturally melds with those of ripples, making detection difficult. When the sonar look-direction is nearly orthogonal to the ripples, ripples produce strong vertical bands of shadow that can be easily rejected with the ripple detector.

VII. DISCUSSION

The standard AUV survey plan used in practice is to execute a series of parallel tracks at one orientation, followed by a second series of tracks at an orientation orthogonal to the first. This “cross-hatching” ensures that each seabed location in the area of interest is observed at least twice. However, the inflexibility of this rigid pre-planned approach introduces severe inefficiencies into the standard sonar data collection procedure.

With a near real-time detection algorithm running onboard the AUV, the survey route can instead be adapted *in situ* based on the data that is collected. The experimental results from Sec. VI provide insight for how this adaptation should be performed.

It was observed in Areas A and C — which were characterized by flat, benign seabed — that a single view was sufficient to detect virtually all objects of interest. Therefore, to save resources, cross-hatching should not be performed in such environments.

To optimize detection performance in areas with sand ripples, rather than blindly executing a second survey in a direction orthogonal to the first, one should instead execute the second survey such that the look-direction of the sonar is orthogonal to the estimated direction of the sand ripples.

The plan for each of the above scenarios — flat seabed or seabed with ripples — relies on an assessment of the environment, which is derived from the results of the ripple detection algorithm. Therefore, the near real-time nature of the detection (and ripple detection) algorithm is the key that would permit this intelligent *in situ* survey-route adaptation.

VIII. CONCLUSION

The overarching idea of the proposed detection algorithm was to overcome the limitations of the most popular existing detection algorithms, while also striving to allow streaming, near real-time detection.

A balance of rigor and flexibility was embedded in the detection algorithm by exploiting definitions (shadows, ripple fields), underlying physics (range-dependent echo loss) and geometry (range-dependent shadow lengths), incorporating extensive domain-specific knowledge (object sizes and signal strengths for template sizes and thresholds), and adapting to *in situ* environmental and data conditions (background reverberation level, sand ripples, image quality).

This work also provides the first quantitative evidence, based on real data or otherwise, that the relative orientation

of sand ripples has a dramatic impact on (non-buried) object detection performance. A suggestion for exploiting this insight to optimize data collection was also provided. A future sea trial will test this proposal using the MUSCLE AUV, to perform real-time detection and survey route adaptation, based on sand ripples, in a completely autonomous manner.

REFERENCES

- [1] B. Walters, J. French, and M. Barnes, “Modeling the effects of crew size and crew fatigue on the control of tactical unmanned aerial vehicles (TUAVs),” in *Simulation Conference Proceedings*, 2000, pp. 920–924.
- [2] S. Williams, O. Pizarro, M. How, D. Mercer, G. Powell, J. Marshall, and R. Hanlon, “Surveying nocturnal cuttlefish camouflage behaviour using an AUV,” in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2009, pp. 214–219.
- [3] Z. Reut, N. Pace, and M. Heaton, “Computer classification of seabeds by sonar,” *Nature*, vol. 314, pp. 426–428, 1985.
- [4] H. Singh, J. Adams, D. Mindell, and B. Foley, “Imaging underwater for archaeology,” *Journal of Field Archaeology*, vol. 27, no. 3, pp. 319–328, 2000.
- [5] Y. Petillot, S. Reed, and J. Bell, “Real time AUV pipeline detection and tracking using side scan sonar and multi-beam echo-sounder,” in *Proc. IEEE OCEANS*, 2002, pp. 217–222.
- [6] Y. Zhang, X. Liao, and L. Carin, “Detection of buried targets via active selection of labeled data: application to sensing subsurface UXO,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, no. 11, pp. 2535–2543, 2004.
- [7] A. Dabbagh, K. Al-Hinai, and M. Khan, “Detection of sand-covered geologic features in the Arabian Peninsula using SIR-C/X-SAR data,” *Remote Sensing of Environment*, vol. 59, no. 2, pp. 375–382, 1997.
- [8] 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.
- [9] 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, 2003.
- [10] 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.
- [11] F. Maussang, J. Chanussot, A. Héret, 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.
- [12] 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.
- [13] P. Viola and M. Jones, “Robust real-time object detection,” *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2004.
- [14] G. Widmer and M. Kubat, “Learning in the presence of concept drift and hidden contexts,” *Machine Learning*, vol. 23, pp. 69–101, 1996.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] 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.
- [19] E. Reffet, S. C. du Pont, P. Hersen, and S. Douady, “Formation and stability of transverse and longitudinal sand dunes,” *Geology*, vol. 38, no. 6, pp. 491–494, 2010.
- [20] A. Lyons, D. Abraham, and S. Johnson, “Modeling the effect of seafloor ripples on synthetic aperture sonar speckle statistics,” *IEEE J. Oceanic Engineering*, vol. 35, no. 2, pp. 242–249, 2010.
- [21] R. Urick, *Principles of Underwater Sound*. McGraw-Hill, 1983.