Copy No. _____

Defence Research and Development Canada
Recherche et développement pour la défense Canada

DEFENCE DÉFENSE

Lessons learned from the DSTO Sonar ATR Challenge

John A. Fawcett
Warren A. Connors

Defence R&D Canada – Atlantic

Technical Memorandum
DRDC Atlantic TM 2013-184
February 2014

Canada

Lessons learned from the DSTO Sonar ATR Challenge

John A. Fawcett
Warren A. Connors

Defence Research and Development Canada – Atlantic
Technical Memorandum
DRDC Atlantic TM 2013-184
February 2014

Abstract

In May 2013 DRDC Atlantic participated in a Defence Science and Technology Organization, Australia (DSTO)-sponsored sonar automatic target recognition (ATR) challenge. As part of this challenge, DRDC Atlantic and the other participants received approximately 11000 Remus/Marine Sonic sonar images in bitmap format as well as corresponding navigational information. A set of 654 sonar files and ground truth data for minelike objects and objects of interest were supplied for training the ATR methods. A larger set of files containing minelike and objects of interest were then used for testing the methods. The locations of the objects in the testing set were not known to the participants. A file with the locations of the DRDC Atlantic ATR detections was emailed to DSTO and the results were evaluated on the basis of the ground truth locations for the testing set. In this report, we describe our initial processing of the data to obtain the results which were sent to Australia. The set of sonar files provided were very challenging, including significant clutter regions and sand ripples. The results from the test evaluation indicated a low detection rate. After the test evaluation, Australia kindly supplied a file with the ground truth locations for the testing set. This report describes the subsequent analysis of our ATR approach. It was found that by dropping one of the constraints originally used in the ATR algorithm that the results were significantly improved. The results were further improved by taking the local clutter density into account. Finally, a binary classifier, trained with shadow/highlight features, improved the detection results even further.

Résumé

En mai 2013, Recherche et développement pour la défense Canada (RDDC) Atlantique a participé à l'épreuve de reconnaissance automatique des objectifs (ATR) par sonar tenu par la Defense Science and Technology Organization (organisation pour les sciences et la technologie pour la défense), d'Australie. Les participants, dont RDDC Atlantique, ont reçu environ 11 000 images sonar Remus-Marine Sonic en format bitmap, ainsi que les données de navigation connexes. Ils ont aussi reçu un ensemble de 654 fichiers sonar comportant des données de réalité de terrain sur des objets d'intérêt et d'autres ressemblant à des mines, afin de leur permettre de régler leurs méthodes d'ATR. Ces méthodes ont ensuite été mises à l'essai à l'aide d'un ensemble plus vaste de fichiers sur d'autres objets d'intérêt ou ressemblant à des mines, dont la position était inconnue des participants. RDDC a soumis par courriel à la DSTO un fichier indiquant la position des objets détectés par ses méthodes d'ATR ; ces résultats ont été évalués d'après les positions de réalité de terrain de l'ensemble d'essai. Dans le présent rapport, nous décrivons tout d'abord le traitement initial des données dont nous avons envoyé les résultats à la DSTO. L'ensemble de fichiers sonar fourni était fort complexe, car il comprenait des zones considérablement

encombrées et des ondulations de sable, et nous avons obtenu de piètres résultats. Après l'évaluation, la DSTO nous a généreusement remis un fichier indiquant les positions réelles des objets à détecter dans l'ensemble d'essai. Le rapport traite aussi de l'analyse subséquente de notre méthode d'ATR. Nous avons ainsi découvert trois façons d'améliorer considérablement la détection des objets : surtout lever l'une des contraintes de l'algorithme ATR, mais aussi tenir compte de la densité du fouillis local et utiliser un classificateur binaire entraîné à l'aide de caractéristiques ombre-lumière.

Executive summary

Lessons learned from the DSTO Sonar ATR Challenge

John A. Fawcett, Warren A. Connors; DRDC Atlantic TM 2013-184; Defence Research and Development Canada – Atlantic; February 2014.

Background: DRDC Atlantic and other organizations, including national research centres, universities, contractors, were invited by Defence Science and Technology Organization, Australia (DSTO) to participate in the "DSTO Sonar ATR Challenge". DRDC Atlantic agreed to participate in this challenge and over approximately one week in May 2013 processed approximately 11000 Marine Sonic sonar images with one of our ATR algorithms. A set of 654 sonar files was provided with ground truth information for algorithm training purposes. The testing set was not provided with ground truth information and hence was a blind test. Each participant provided DSTO with a file of pixel locations for their ATR detections and the results were evaluated by DSTO with the ground truth information. The results of the test evaluation for the DRDC Atlantic ATR method were somewhat disappointing. The detection rate with the testing set was low with only 39% of the minelike objects detected. After the evaluation, DSTO provided DRDC Altantic with the testing set ground truth allowing for an analysis of the ATR algorithm's performance. This analysis, which is described in this report, indicated that our performance was very significantly improved by removing a constraint on a matched filter output which had been used in combination with the output from a Haar-cascade face detector.

Principal results: Automatic target recognition (ATR) methods can be trained and the various algorithm parameters varied to yield optimal performance for a particular sonar environment and sonar. It is much more challenging to obtain good ATR performance in new environments possibly with targets not previously seen. The DSTO sonar ATR challenge data set provided an excellent opportunity to test an ATR method on a difficult and varied data set without much time to fine tune the algorithms. Subsequent analysis showed that the results could be significantly improved by changing the constraints and thresholds used within the ATR algorithm. In addition, a secondary classification stage can further improve the results. Although, these improved results were encouraging, this post-analysis emphasizes the difficulty in choosing, a priori, good parameter values for the ATR algorithms.

Significance of results: It is shown that on a large and difficult set of Remus/Marine Sonic data, it was possible to obtain good ATR results. However, this does require some care in choosing the internal threshold values carefully. In general, the ATR needs to be calibrated when entering into a new sonar environment either by human or computer analysis. The use of large and varied sonar data sets (with groundtruth information) as blind test sets is a very useful exercise in testing ATR methods.

Future work: The automatic adaptation of ATR methods to new sonar environments will be investigated. As the DSTO sonar ATR challenge proved to be very worthwhile, it is recommended that DRDC participate in similar challenges in the future.

Sommaire

Lessons learned from the DSTO Sonar ATR Challenge

John A. Fawcett, Warren A. Connors ; DRDC Atlantic TM 2013-184 ; Recherche et développement pour la défense Canada – Atlantique ; février 2014.

Contexte : Defence Science and Technology (DSTO, ou organisation pour les sciences et la technologie pour la défense), d'Australie, a invité Recherche et développement pour la défense (RDDC) Atlantique, des centres nationaux de recherches, des universités et des entrepreneurs, entre autres, à participer au DSTO Sonar ATR Challenge (épreuve de reconnaissance automatique des objectifs [ATR] par sonar de la DSTO). RDDC Atlantique a accepté l'invitation et a traité en mai 2013 durant une semaine environ approximativement 11 000 images sonar Marine Sonic au moyen d'un de ses algorithmes d'ATR. On nous a remis un ensemble de 654 fichiers sonar comportant des données de réalité de terrain, afin de régler nos algorithmes. L'épreuve, cependant, consistait en un essai aveugle, car l'ensemble d'essai ne comportait aucune donnée de réalité de terrain. Chaque participant a soumis à la DSTO un fichier de positions des pixels de ses détections ATR, puis celle-ci a évalué les résultats daprès les données de réalité de terrain. Nous avons obtenu des résultats plutôt décevants, car notre algorithme dATR n'a détecté que 39 % des objets ressemblant à des mines. Après l'évaluation, DSTO nous a donné les positions réelles des objets à détecter, afin que nous puissions analyser le rendement de notre algorithme. Cette analyse, décrite dans le présent rapport, a montré que l'on pouvait considérablement améliorer la détection en levant la contrainte sur la sortie de filtre adapté utilisée de pair avec la sortie d'un détecteur de visages par cascade de Haar.

Résultats : Nous pouvons obtenir un rendement optimal dans un environnement sonar et pour un sonar particulier en réglant les méthodes d'ATR et en faisant varier les divers paramètres des algorithmes, mais détecter automatiquement les cibles dans de nouveaux environnements comportant potentiellement des objectifs non détectés auparavant est plus ardu. L'ensemble de données du DSTO Sonar ATR Challenge a été une excellente occasion de mettre à l'essai nos algorithmes d'ATR sur un ensemble de données complexe et varié sans avoir le loisir de les régler finement. Une analyse subséquente a montré que les résultats pourraient être grandement améliorés en modifiant les contraintes et les seuils de l'algorithme d'ATR et en ajoutant une étape de classification secondaire. Ces résultats améliorés sont certes encourageants, mais l'analyse après-essai souligne la difficulté de régler adéquatement et a priori les paramètres des algorithmes d'ATR.

Portée : L'essai a montré qu'il est possible d'obtenir de bons résultats d'ATR sur un vaste et complexe ensemble de données Remus et Marine Sonic, mais que cela exige

de choisir soigneusement les valeurs de seuil internes. Il faut généralement étalonner l'ATR dès l'entrée dans un nouvel environnement sonar, au moyen d'une analyse manuelle ou informatique. Faire l'essai aveugle de nos algorithmes d'ATR sur des ensembles de données sonar lourds et variés, doublés des données de réalité de terrain pour les valider, s'est montré fort utile.

Recherches futures : L'adaptation automatique de méthodes ATR à de nouveaux environnements sonar sera examinée. Étant donné que la participation de RDDC à l'épreuve ATR par sonar de la DSTO fut jugée très profitable, il est recommandé de répéter ce genre d'expérience à l'avenir.

Table of contents

List of figures

List of tables

Acknowledgements

We would like to thank DSTO Australia for organizing the Sonar ATR Challenge and for providing us with the sonar images, and after our participation, the groundtruth information.

This page intentionally left blank.

1 INTRODUCTION

In May 2013, the Defence Science and Technology Organisation (DSTO) Australia distributed a large data set of REMUS MarineSonic images that DSTO had collected over the years. Some of these images contained minelike objects. A total of 654 bitmap images were supplied as a training set along with an Excel file giving the file/pixel locations of minelike objects and objects of interest. A total of 11202 bitmap images comprised the testing set. In the case of the testing set, the groundtruth information was not provided during the challenge. The testing images included a wide variety of seabed and sonar conditions, including images with very significant clutter. Four representative images from the testing set are shown in Fig. 1. Figures 1a and 1c show seabeds with significant clutter. The images of Figs. 1b and 1d are more benign although Fig. 1d shows a significant scour mark. In Figs. 1a and 1d there are artifacts caused by Autonomous Underwater Vehicle (AUV) turns in the images. There is also significant altitude variation in Fig. 1a. In all images, a surface echo can be seen with varying degrees of amplitude.

The training set was used to train a Haar-cascade[1,2] with the openCV software [3]. This cascade was then used along with a matched-filter[4,5] to detect objects in the testing set. In order to do this, we had to make a choice as to what thresholds to use in terms of the number of cascade-detection rectangles and the matched-filter output associated with the pixels in the image. We chose these values by trying the method on a few images of the testing set and observing the resulting detections. Due to the large number of files, the processing procedure was quite lengthy and it was divided between the two authors of this report. We performed the training and testing over a period of about a week and did not have much time to fine tune our methods. Also, due to the time constraint, we did not implement a secondary classification procedure to further reduce the number of false alarms.

A text file of all detections, consisting of their x/y pixel locations for each file, was emailed to DSTO who then evaluated the detector performance on the basis of their groundtruth information. The evaluation was done by using a distance threshold between the detector's pixel locations and the groundtruth locations. The scoring was somewhat complicated. There were groundtruth locations of minelike objects (MLO) marked by Australian naval personnel. In addition, there were files for which a human detection was not made but in which the minelike object should appear (according to the geographical location of the sonar file). There were also designated objects-of-interest in the images. From the evaluation, the DRDC results were somewhat disappointing, the detection rate of the basic 388 MLOs was only 39%, although there were 54 additional MLO detections in the other files.

It was clear from the low detection rate in the testing set that the values of the thresholds we used were not appropriate. Since the evaluation, DSTO Australia has

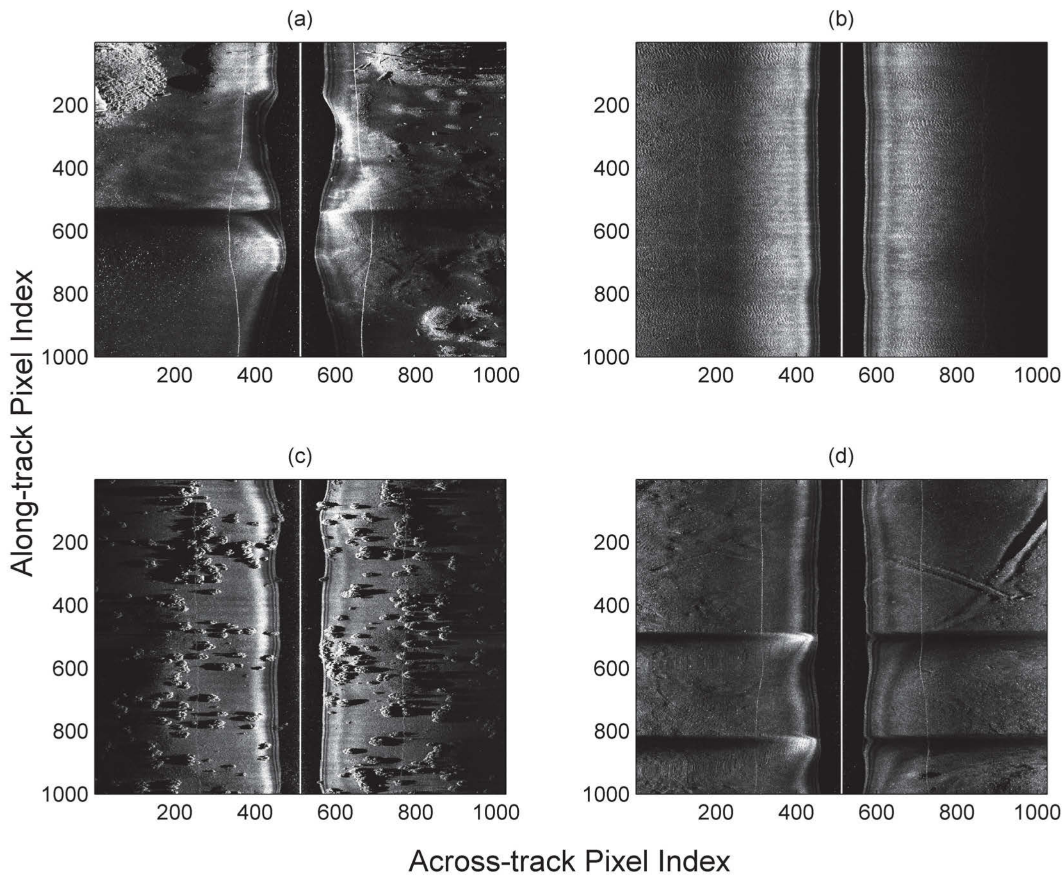

Figure 1: *Four representative images from the DSTO Australia set.*

provided DRDC Atlantic with the ground-truth locations allowing for the investigation of ATR algorithm's performance. In addition, a secondary classification process, derived from the training set, was implemented to further reduce the number of false alarms. Below, we describe the algorithms and the results.

2 DETECTORS

2.1 Haar Cascade Detectors

The Haar-Cascade detector is commonly used in face detection in images and has recently been applied to sidescan sonar ATR [1,2]. In this approach, windows of the image are passed through a series of increasingly complex detector stages. Each stage is trained to have a very high detection rate but to significantly eliminate false alarms. If a window is rejected at an initial stage, it is not considered by the subsequent stages. In a Haar-Cascade, the individual detectors utilize Haar-features which can be rapidly computed. We utilize the cacscade training software in the well-known openCV [3] software library. As well, we use the trained cascade with the multi-scale face detection algorithm from openCV for the detection phase. This algorithm has been modified to save its detection rectangles into a text file which can be read into a MATLAB program.

In this MATLAB program the centres of all the rectangles are computed for each file. A counting matrix C is constructed where each (i, j) location has a 1 added to it if it corresponds to a rectangle centre. In other words, the value at each pixel location is the number of rectangle centres at that point. This counting matrix is then convolved with a 30×30 filter of ones to provide a smoothed count matrix, C_S, of the rectangles. The number of rectangles required for a detection is a parameter which can be varied to generate a Receiver Operating Characteristic (ROC) curve. The dimensions of the rectangles that are counted can be constrained. For the results of this report, the width (across-track) of the rectangles were such that $16 \leq w \leq 100$ and the heights (along-track) satisfied $12 \leq h \leq 30$ pixels. We also computed a normalized rectangle count. This is done by convolving C_S with a filter which is 200×150 in size, with the middle 31×31 pixels set to zero to compute an averaged value C_m. Then the normalized rectangle count is given by

$$C_n = \frac{C_S}{1 + C_m}. \tag{1}$$

In training a Haar Cascade, a set of "positive" images and a large set of negative images is required. The positive images are cropped images of minelike objects while the negative images are simply sonar images without minelike objects. The bitmap images were read into a MATLAB program and the images were normalized - that is, the overall range/amplitude variation was empirically removed. This was done by computing the mean amplitudes along predicted constant grazing angles from the sonar. To obtain the images to crop for the positive images, we used the ground-truth locations provided by DSTO for the training set, to locate the minelike objects and also the objects of interest. From the predicted locations in the sonar files, we used a mouse to extract a rectangular region about both the minelike and objects of

Parameter	Value
Window Width	16
Window Height	12
Number of Stages	17
Hit Rate	0.999
Max False Alarm Rate	0.4
Haar Feat. Set	ALL
Maximum Depth	2

Table 1: *The parameters used for training the cascade.*

interest. The number of positive images was increased by first considering the basic extracted snippet of data and then perturbing the location of the extracted rectangle and also by distorting the width of the original snippet. In addition, we utilize both the normalized and unnormalized versions of the images. The maximum width of an extracted rectangle is 40 pixels, meaning that long shadows are truncated. In Fig. 2 we show a portion of the set of positive snippets. There was a fairly wide variety of minelike and objects of interest. There were standard target types but also some smaller objects and slender cylinders. The many images which did not contain the indicated objects were taken as negative (clutter) files. We used a total of 4528 clutter images (in the TIFF format). The starboard and port sides of a sonar file are treated as separate images. In addition, normalized and unnormalized versions, and vertically flipped and non-flipped variants of the original images are used. The openCV software was then used to train a Haar Cascade. The parameters of this cascade are listed in Table I.

2.2 Matched Filtering

One aspect of a sidescan sonar image that is not accounted for in the cascade training is the fact that a object's shadow length grows linearly, in terms of the number of across-track pixels, as a function of detection range. It might be possible to incorporate this in the Haar features but we have not attempted to modify the openCV software for this task. However, the standard matched filtering used in sidescan or Synthetic Aperture Sonar (SAS) ATR [4,5] is, in fact, a generalized type of Haar filter. The concept is to construct a two-dimensional filter T with a small rectangle of ones to model the highlight section of an object followed by a wider (in across-track direction) rectangle of -1 to model the shadow. Let us consider an object with height h, let us take the altitude of the sonar as a and the range (slant) of the detection is R. Taking all units to be in metres, the length of the shadow is given approximately

Figure 2: *Some representative positive images from the training set.*

by

$$L_S = \frac{hR}{(a-h)}, \tag{2}$$

which can be converted into a pixel length according to the across-track sampling. This length will define the extent of the shadow region for the template T. The original image I is remapped, either in discrete or continuous fashion, into a new image I_m having positive and negative values. For the challenge, the following remapping was used: (1) first the median ν of the image I pixel values (outside of the water column) was determined and then (2) the remapped image

$$I_L = \frac{\ln(\frac{I+.001}{\nu})}{\ln(2)} \tag{3}$$

was computed. The mapping of Eq.(3) corresponds to image pixel values with amplitudes twice the median value being mapped into the value 1, values 4 times the median value mapped into 2, values 1/2 the median value mapped into -1, etc. Finally, for the remapped image the negative values of I_L are rescaled by 1/3 and clipped at -1. Then the two-dimensional cross-correlation of I_L with T will have a high output when there is a good match between the remapped image and T. Using the method of integral images [6], it is straightforward to allow the final cross-track pixel indices of the shadow rectangle to increase as the range of the leading pixel increases. In this report, we will consider a matched filter with the highlight and shadow sections independently normalized, thus weighting the highlight and shadow equally in the output value.

2.3 Feature-based Classifier

For the original detection results sent to DSTO, only the combination of a Cascade and a matched-filter were used. However, for later analysis a trained kernel-regression method[7] was also used to process the data snippets (mugshots) yielded by the initial detection phase on the training data set. From the training set ground truth, the mugshots are assigned a label as either a MLO or non-MLO. This classifier uses several computed shadow/highlight features. Although we have called this method a classifier, it is being used simply to further classify mugshots or snippets of data as minelike or clutter. In this sense, it is still considered as a detection stage. The mugshots are first segmented[8] into a shadow and highlight region. A number of features based upon these regions, such as estimated height of the object, the estimated length, the area of the shadow , the ratio of this area to the area of a fit ellipse, are computed. These features are described in [9]. Additional features from the detection phase are also included: (1) the number of associated rectangles (2) the matched-filter output with the original detection (3) the normalized rectangle count and the maximum *lacunarity* [4] in the detection region. The labelled set of features,

$\{\vec{f_i}\}$, are then used to train a Kernel-regression classifier,

$$K(i, j) \equiv exp(-|\vec{f_i} - \vec{f_j}|/p) \tag{4}$$

and

$$K\vec{a} = \vec{\ell} \tag{5}$$

where $\vec{\ell}$ denotes the labels and p is a parameter which can be varied and effectively controls the width of the exponential term of the kernel. A small scalar may be added to the diagonal of K in Eq.(5) to regularize the solutions.

Once the coefficient vector $\vec{a}$ has been determined from the training, then for a new feature $\vec{f_m}$, the predicted label value is given by

$$\ell_m = \sum_{i=1}^{N} a_i K(\vec{f_m}, \vec{f_i}). \tag{6}$$

3 RESULTS

For the results that were sent to DSTO, it was required that there be at least 6 rectangles associated with a detection and that it have a match-filtered output of at least one. Due to the mapping of the image data into positive and negative values, and the manner in which the matched-filter template is normalized, the matched-filter values can often exceed unity. We ignored any detections which were within a region designated as a turn or were in the water colum. The turn and water column regions were automatically determined from the navigational data. For the results sent to DSTO, there was a bug in the code for determining the detections within the water column. This bug was corrected for the results of this report. Two text files were constructed with the detection information, one including all the detections and their pixel locations within the sonar image files and the second file was the same as the first but with the number of detections in a file limited to a maximum of 5 for each side (port and starboard). The five detections with the highest rectangle count were used. In this latter case, we achieved about a 39% detection rate on the the MLO with about a total of 7000 false alarms.

For the results of this report, the initial detections corresponded to those regions where the number of rectangles are greater than or are equal to two. Within these regions, the maximum values of the rectangle count and the other detectors, matched filter, lacunarity, and a normalized version of the rectangle, Eq.(1) are saved as well as the pixel indices of the detection. The detection is declared a true detection if it is within a distance of 30 pixels from the ground truth location. Here, the ground truth detections used are only the designated MLOs in the ground truth file. By varying the threshold of the detectors' outputs ROC curves are generated. In Fig. 3 the ROC curve obtained by varying the threshold on the rectangle count is shown (blue). The same curve when a constraint (≥ 1) upon the matched-filter output is also imposed is shown (red). Finally, the curve generated by varying the threshold upon the normalized rectangle count is shown. The red curve corresponds, approximately, to the parameters used for the results sent to Australia and it can be seen that it was the constraint upon the matched-filter output which caused the low detection rates. The value of unity as a constraint for the match-filtered output was too high and led to a low detection rate. We can lower this threshold to obtain better detection rates, but in general, for this particular data set, the matched filter output was not useful in improving the detection results from the cascade detector. This may be due to the fact that many of the false alarms were produced by rocks, coral, etc for which the matched-filter output is relatively high. For the detection results initially sent to DSTO, we had a rectangle threshold of 6, which on the unconstrained rectangle-count curve corresponds to a detection rate of 85% and 40448 false alarms. Thus, even though the ROC curve indicates good performance for the cascade approach, the particular threshold we used initially was too low. A more appropriate threshold would have been 20 for which the probability of detection is 70% and there are 11555 false alarms. The normalized rectangle count was the best detector in this case, yielding a 70% detection rate with approximately 5000 total false alarms.

Some files produced a large number of detections (false alarms). In practice, one might simply designate those files as "clutter" files with the corresponding geographical locations declared as "non-huntable". Another strategy is to limit the number of detections per file. In this analysis the port and starboard sides are treated as separate files . The detections retained are those with the highest detector scores. In Fig.4, we consider the unnormalized rectangle score with no constraint on the number of detections/file (blue), 8 detections/file(red), 4 detections/file (green), 2 detections/file (cyan), and finally one detection/file (yellow). From Fig. 4 it can be seen that limiting the number of detections per file has had a beneficial effect. The limit of only one detection/side would be, in general, too restrictive - however values of 4 and 8 can also be seen to be useful. In Fig. 5 the corresponding curves are shown when the normalized rectangle score is used. In this case, the improvement is less significant, although at the higher rates of detection the improvement is noticeable. Using a normalized score effectively adjusts the threshold in terms of the number of nearby false alarms and hence already accomplishes the concept of accounting for the number of detections within a file.

Finally, we consider the mugshots resulting from the detections for the training set. These detections were obtained using a rectangle threshold of 2. The corresponding mugshots are then read into a MATLAB segmentation/feature computation program. As well, from the detection phase, the values of the rectangle count, the normalized rectangle count, the matched-filter output, and the lacunarity are included as features. In total, there are 27 features for each mugshot. A kernel regression method is trained with the parameter p in Eq.(4) equal to 10. The computed coefficients $\vec{a}$ are then used to compute output values, Eq.(6), for the features of the testing set and it is a threshold on this output value which is varied to produce the ROC curve shown in Fig. 6. As a comparison, the ROC curve for the normalized rectangle count is shown in red. The upper plot shows the entire ROC curves while the bottom plot shows a zoom of the 2 ROC curves for lower false alarm rates. In the upper plot, the 2 curves appear quite similar but the bottom figure does show that, for example, at a 70% detection level, the kernel regression method has approximately 3500 false alarms while the normalized rectangle count has approximately 4500 false alarms. Thus, the trained kernel method outperforms the normalized cascade method but not greatly.

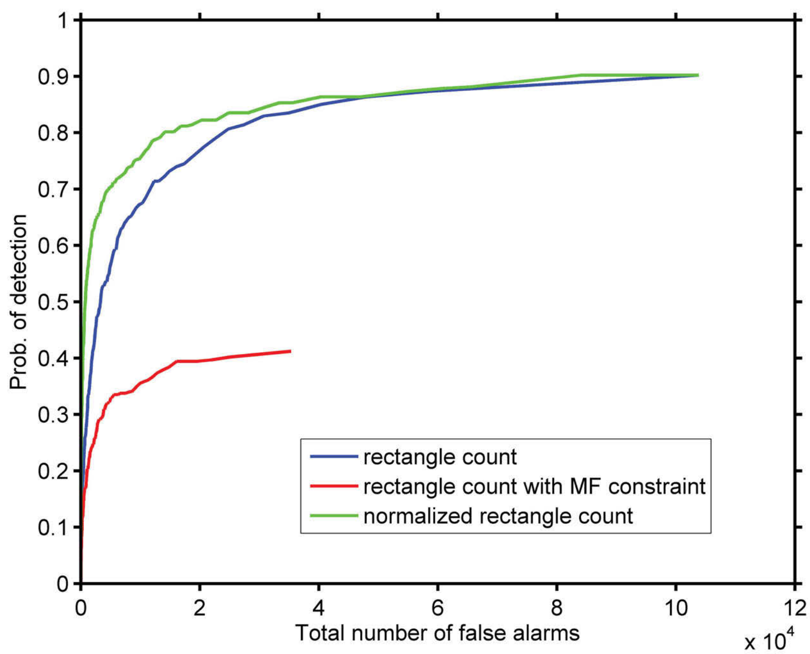

Figure 3: *The ROC curves for the testing set using the rectangle count from the Haar-cascade:(a) unnomalized count (blue) (b) unnormalized count with constraint on match-filter output (red) (c) normalized count (green)*

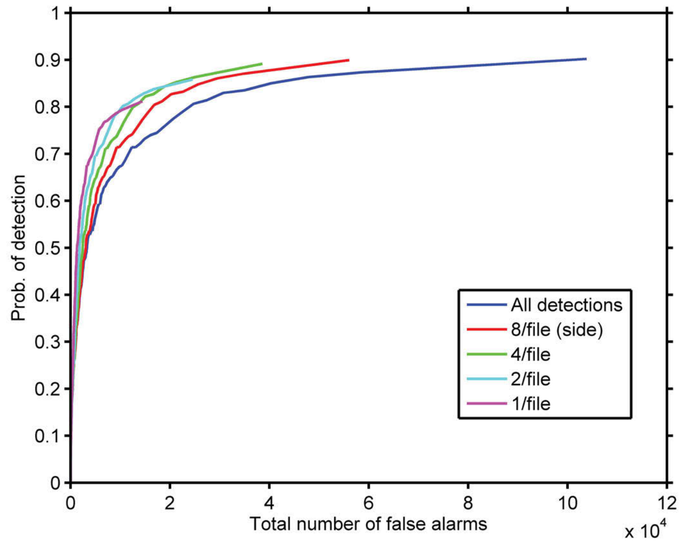

Figure 4: *The ROC curves for the testing set using the unnormalized rectangle count from the match-filter output and with different numbers of limits on the maximum number of detections per side*

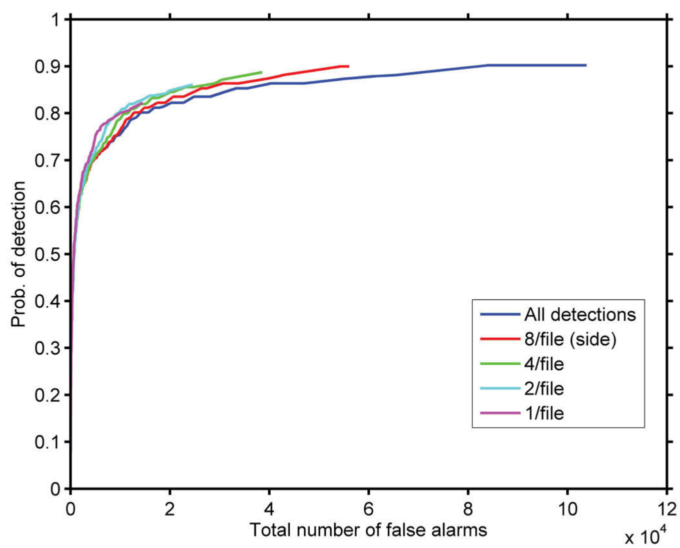

Figure 5: *The ROC curves for the testing set using the normalized rectangle count from the match-filter output and with different numbers of limits on the maximum number of detections per side*

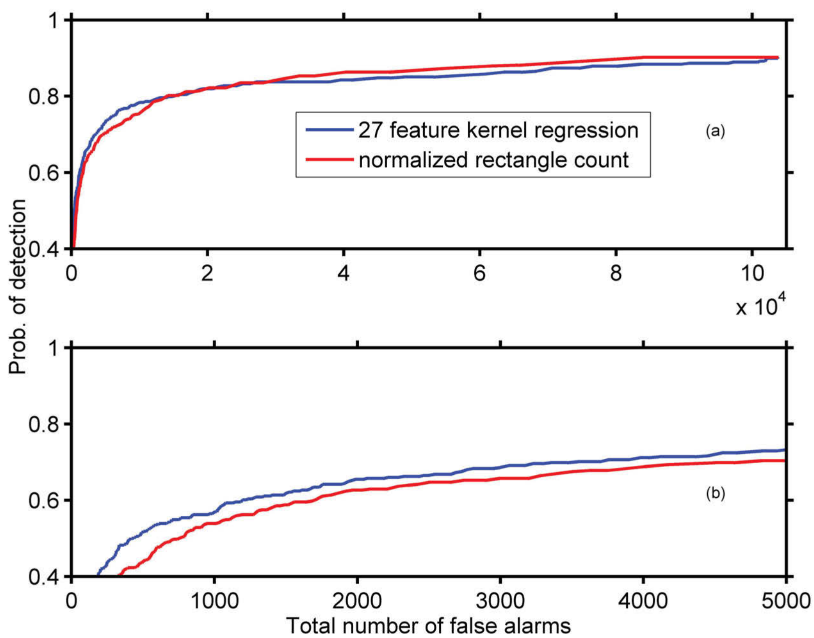

Figure 6: *The ROC curves for the testing set using the output from a trained kernel-regression method (blue) and for the normalized rectangle count (red) . In (a) the entire curves are shown and in (b) a zoom of the lower false alarm region.*

4 DISCUSSION OF RESULTS

The Australian DSTO Sonar Challenge proved to be very interesting. The collection of sonar images proved to challenging from an ATR perspective. There was a wide variety of different seabed conditions and some of the targets were quite thin or small. For a simple detection approach such as the matched-filter this resulted in many false alarms for a high detection rate. The Challenge also highlighted the problem of not knowing, a priori, optimal values to use for thresholds. Without having much time to “fine-tune” ATR parameters, we quickly set what we considered to be reasonable values for the rectangle and matched-filter thresholds. The analysis of this report indicates that our constraint upon the matched filter output was too high and this resulted in a low detection rate. In fact, the matched-filter output was not a particularly useful feature for this data set. The rectangle count from the Haar-feature face detection program performed very well on its own. A normalized rectangular count performed even better than the unnormalized count. This is probably due to the fact that this approach effectively adapts its threshold with respect to the amount of surrounding clutter. However, even though the ROC curves for the Haar-detection rectangles indicated good ATR performance, it is difficult to know at the beginning of processing which particular threshold value to use. For our initial results sent to DSTO, the threshold value of 6 that we used was, in fact, suboptimal and a higher value would have been better. It was also found that applying a hard constraint upon the number of detections per side improved the ROC curves. This approach relies on assuming that the number of targets expected in a single sonar image will be below some specified number. Finally, it was found that a trained kernel-regression method, using the detection outputs and features computed for shadow/highlight regions, yielded the best results. For the testing set of 388 MLOs we obtained about 3500 false alarms for a detection rate of 70% and about 14000 false alarms for a detection rate of 80% with the normalized rectangle count.

References

[1] J.Sawas, Y.Petillot, and Y.Pailhas, “Cascade of boosted classifiers for rapid detection of underwater objects,” in *Proceedings of ECUA 2010*, (Istanbul, Turkey), 2010.

[2] Y.Petillot, Y.Pailhas, J.Sawas, N.Valeyrie, and J.Bell, “Target recognition in synthetic aperture sonar and high resolution side scan sonar using auvs,” in *Proceedings of International Conference: Synthetic Aperture Sonar and Synthetic Aperture Radar, Institute of Acoustics Proceedings*, (Lerici, Italy), Sept. 2010.

[3] Open Source Computer Vision OpenCV 2.4.3 , http://opencv.willowgarage.com/wiki/, Access date:(January 2013).

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

[5] J. Fawcett, A. Crawford, D. Hopkin, V. Myers and B.Zerr, "Computer-aided detection of targets from the CITADEL trial Klein sonar data", DRDC Atlantic TM 2006-115, November 2006.

[6] D. Williams, "On adaptive underwater object detection," in *Proceedings of IEEE/RSJ Interntaional Conference on Intelligent Robots and Systems (IROS)*, (San Francisco, U.S.A), pp. 4741–4748, Sept. 2011.

[7] J. Shawe-Taylor and N. Cristianini, *Kernel Methods for Pattern Analysis*, Cambridge University Press, Cambridge, 2004.

[8] V. Myers, "Image segmentation using iteraion and fuzzy logic" *in Proceedings of CAD/CAC Conference*, Halifax, N.S., 2001.

[9] J.Fawcett, V. Myers, D. Hopkin, A. Crawford, M. Couillard, B. Zerr, "Multiaspect Classification of Sidescan Sonar Images: Four Different Approaches to Fusing Single-Aspect Information", *IEEE Journal of Oceanic Engineering*, Vol.35, pp.863–876, Oct. 2010.

This page intentionally left blank.

DOCUMENT CONTROL DATA

(Security markings for the title, abstract and indexing annotation must be entered when the document is Classified or Designated.)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.) Defence Research and Development Canada – Atlantic PO Box 1012, Dartmouth NS B2Y 3Z7, Canada	2a. SECURITY MARKING (Overall security marking of the document, including supplemental markings if applicable.) UNCLASSIFIED 2b. CONTROLLED GOODS (NON-CONTROLLED GOODS) DMC A REVIEW: GCEC APRIL 2011

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)

Lessons learned from the DSTO Sonar ATR Challenge

4. AUTHORS (Last name, followed by initials – ranks, titles, etc. not to be used.)

Fawcett, J. A.; Connors, W. A.

5. DATE OF PUBLICATION (Month and year of publication of document.) February 2014	6a. NO. OF PAGES (Total containing information. Include Annexes, Appendices, etc.) 30	6b. NO. OF REFS (Total cited in document.) 9

7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Technical Memorandum

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

Defence Research and Development Canada – Atlantic
PO Box 1012, Dartmouth NS B2Y 3Z7, Canada

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.) 11cc	9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)
10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.) DRDC Atlantic TM 2013-184	10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

(X) Unlimited distribution
(   ) Defence departments and defence contractors; further distribution only as approved
(   ) Defence departments and Canadian defence contractors; further distribution only as approved
(   ) Government departments and agencies; further distribution only as approved
(   ) Defence departments; further distribution only as approved
(   ) Other (please specify):

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11)) is possible, a wider announcement audience may be selected.)

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

In May 2013 DRDC Atlantic participated in a Defence Science and Technology Organization, Australia (DSTO)-sponsored sonar automatic target recognition (ATR) challenge. As part of this challenge, DRDC Atlantic and the other participants received approximately 11000 Remus/Marine Sonic sonar images in bitmap format as well as corresponding navigational information. A set of 654 sonar files and ground truth data for minelike objects and objects of interest were supplied for training the ATR methods. A larger set of files containing minelike and objects of interest were then used for testing the methods. The locations of the objects in the testing set were not known to the participants. A file with the locations of the DRDC Atlantic ATR detections was emailed to DSTO and the results were evaluated on the basis of the ground truth locations for the testing set. In this report, we describe our initial processing of the data to obtain the results which were sent to Australia. The set of sonar files provided were very challenging, including significant clutter regions and sand ripples. The results from the test evaluation indicated a low detection rate. After the test evaluation, Australia kindly supplied a file with the ground truth locations for the testing set. This report describes the subsequent analysis of our ATR approach. It was found that by dropping one of the constraints originally used in the ATR algorithm that the results were significantly improved. The results were further improved by taking the local clutter density into account. Finally, a binary classifier, trained with shadow/highlight features, improved the detection results even further.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

ATR, Remus, sonar, Haar cascade

This page intentionally left blank.

Defence R&D Canada

Canada's leader in defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale

DEFENCE DÉFENSE

www.drdc-rddc.gc.ca