Real-Time Performance of Fusion Algorithms for Computer Aided Detection and Classification of Bottom Mines in the Littoral Environment

Charles M. Ciany & William C. Zurawski
Raytheon Company
Integrated Defense Systems
1847 West Main Road
Portsmouth, RI 02871-1087

Gerald J. Dobeck & Dennis R. Weilert
NSWC Coastal Systems Station
Dahlgren Division
Highway 98
Panama City, FL 32407-7001

***Abstract*- The fusion of multiple Computer Aided Detection/Computer Aided Classification (CAD/CAC) algorithms has been shown to be effective in reducing the false alarm rate associated with the automated classification of bottom mine-like objects when applied to side-scan sonar images taken in the littoral environment. Real-time operation of the CAD/CAC fusion algorithms from Raytheon, Lockheed Martin, and NSWC Coastal Systems Station (CSS) on board an unmanned underwater vehicle has recently been successfully demonstrated as part of a littoral test sponsored by the Office of Naval Research (ONR) in 2002. Test results proved that the fusion reliably classified bottom mine-like objects while significantly reducing the false alarm rate relative to that of a single CAD/CAC algorithm. This paper discusses the hardware and software architecture for the real-time implementation of the CAD/CAC algorithms, and presents the real-time performance results obtained during the experiment. Additional post processing performance results are also discussed for alternate fusion approaches, and the overall performance benefit through a significant reduction of false alarms at high correct classification probabilities is quantified.**

I. CAD/CAC FUSION OVERVIEW

Over the past several years, CAD/CAC processing algorithms from Raytheon [1-3], Coastal Systems Station (CSS) [9,10], and Lockheed Martin (LM) [11] have been adapted to the littoral environment by processing side-scan sonar images recorded by the REMUS autonomous underwater vehicle (AUV) to demonstrate reliable classification of bottom mines. The combination of the outputs from the multiple CAD/CAC algorithms, termed Algorithm Fusion, has been applied in non-real time by post-processing the recorded sonar images, and the results shown to be beneficial in reducing the overall false alarm rate [2-4,5-12]. The CAD/CAC Fusion processing architecture was recently modified to operate in real-time on-board the REMUS vehicle, and the real-time operation was demonstrated during a sea test sponsored by ONR in 2002.

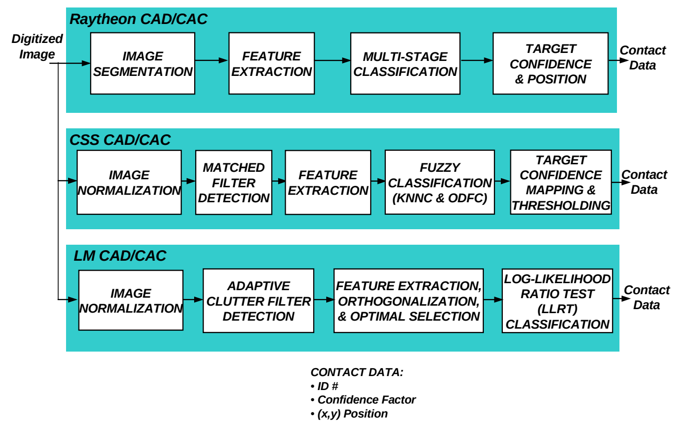

Fig. 1. CAD/CAC processing flow

The CAD/CAC processing algorithms are functionally illustrated in Figure 1. Each algorithm generates output contact data consisting of the contact's planar image coordinates and an associated confidence factor. The normalized confidence factor is a measure of the mine-likeness of the classified contacts, such that valid detections tend to have high confidence factors (approaching 1.0), while false alarms tend to have low confidence factors (approaching 0.5). The Raytheon CAD/CAC processing consists of four primary functions:

- ***Image Segmentation***, in which a subset of the digitized sidescan sonar image is filtered to reduce speckle and then split into non-overlapping segments ("subframes"), each of which is adaptively thresholded (via histograms) to identify Highlight, Shadow, and Background (H/S/B) pixel types;

- ***Feature Extraction,*** in which the H/S pixels of each image segment are geometrically associated to form contiguous H/S regions of interest, each of which are processed to extract key Signal-to-Noise Ratio (SNR) and shape features;
- ***Multi-Stage Classification*** of each region as mine-like or non-mine-like through the association of the individ-

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **01 SEP 2003**	2. REPORT TYPE **N/A**	3. DATES COVERED **-**

4. TITLE AND SUBTITLE **Real-Time Performance of Fusion Algorithms for Computer Aided Detection and Classification of Bottom Mines in the Littoral Environment**	5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Raytheon Company Integrated Defense Systems 1847 West Main Road Portsmouth, RI 02871-1087**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR’S ACRONYM(S) 11. SPONSOR/MONITOR’S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES

See also ADM002146. Oceans 2003 MTS/IEEE Conference, held in San Diego, California on September 22-26, 2003. U.S. Government or Federal Purpose Rights License, The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**UU**	**7**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

ual H/S regions, a weighted scoring of the associated features, thresholding of the weighted scores, and geometric clustering of the regions that pass the threshold; and

- ***Target Confidence & Position Estimation,*** in which a normalized target confidence level and relative location within the image are generated for each classified contact, and output to the subsequent data fusion processor.

The CSS CAD/CAC processing in Fig. 1 consists of the following functions:

- ***Image Normalization,*** in which the sonar image is normalized by the background levels using a dual-pass algorithm, to accentuate target echo and shadow features and remove variations in intensity level associated with different bottom environments;

- ***Matched Filter Detection*** of potential targets of interest using a non-linear filter, matched to the salient characteristics of mine-like objects (MLOs);

- ***Feature Extraction,*** in which a set of candidate target features is extracted from the detected locations of the image;

- ***Fuzzy Classification*** of each detected target to generate a measure of its mine-likeness through processing of the candidate target features with both a K-Nearest Neighbor Attractor-based Neural Network Classifier (KNNC) and an Optimal Discrimination Filter Classifier (ODFC), followed by a fuzzy-logic and-ing of the two classifier outputs [7,8]; and

- ***Target Confidence Mapping & Thresholding,*** in which each detected target is tagged as mine-like or non-mine-like by mapping its mine-likeness measure to a normalized target confidence level and thresholding the result.

The Lockheed Martin CAD/CAC processing of Fig. 1 consists of the following functions:

- ***Image Normalization,*** in which the sonar image is normalized by the background levels using the same dual-pass algorithm as used in the CSS classifier;

- ***Adaptive Clutter Filter Detection*** of potential targets of interest using an adaptive clutter filter tuned over a training set of data to maximize the output response to the average target signature while minimizing the response to non-target signatures;

- ***Feature Extraction & Orthogonalization,*** in which a set of candidate target features is extracted from the detected locations of the image and the features are projected onto orthogonal space;

- ***Log-Likelihood Ratio Test (LLRT) Classification*** of each detected target to generate a measure of its mine-likeness through processing of the candidate target features with a Bayesian classifier that employs stored histograms of target and non-target distributions from previous training to perform a log-likelihood ratio discrimination test; and

The Algorithm Fusion receives the output contact data from each CAD/CAC algorithm and produces fused classification reports of the MLOs. The fusion processing consists of two principal functions:

- ***Geometric Clustering*** of the classified contacts from the multiple classifier outputs, which groups nearby, individually classified contacts together when the distance between them is less than a prescribed threshold; and

- ***Cluster Thresholding,*** in which the target confidence levels associated with clustered contacts are processed and then compared to a threshold for final classification. Several techniques have been considered in previous demonstration studies [13]; the subset employed in this study consists of:

  - ❑ ***Cluster Confidence Factor Sum & Threshold***: the algorithm employed in the real-time fusion processing, in which the average of the confidence levels from the three algorithms must exceed a threshold level, and

  - ❑ ***M-of-N:*** a simple algorithm utilized in the post analysis, which requires at least M=2 of the N=3 classifiers to classify a clustered contact.

II. CAD/CAC FUSION PROCESSING ARCHITECTURE

The real-time CAD/CAC processing flow implemented on board the REMUS AUV is depicted in Figure 2. As the vehicle maps the survey area with its side-scan sonars, the vehicle computer continuously formats the sonar output into image files. As the files are generated, they are saved to disk for later downloading upon retrieval of the vehicle, to support operator post-processing.

After generating each image file, the vehicle computer sends a message to an FTP Client process executing on the embedded CAD/CAC hardware, indicating that the image is ready for processing. Upon receipt of this message, the FTP Client process executes a script that writes the image file to a RAM disk and sends a "ready" message to the real-time

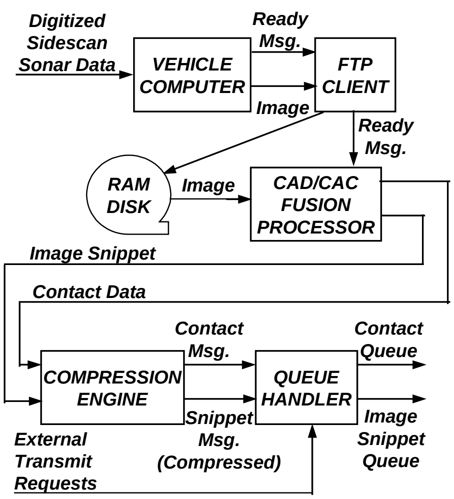

Fig. 2. Real-Time CAD/CAC Fusion Processing Architecture

CAD/CAD Fusion processor, commanding it to process the newly stored image file. The CAD/CAC Fusion process then retrieves the side-scan image from the RAM disk and executes each of the CAD/CAC algorithms to operate on the image (port side first, followed by starboard side). The results from each of the CAD/CAC algorithms are then filtered to exclude designated intervals, such as those in which the vehicle is in an excessive turn.

The filtered contacts from each of the three algorithms are geometrically clustered as described above, and their individual confidence factors are summed. The CAD/CAC processor also extracts a "snippet" of the image from a local region about each of the most mine-like contact clusters whose summed confidence factor exceeds a prescribed threshold. The image snippet and the contact data for each of these clusters are then sent to a Compression Engine processor. Contact data for clusters that don't pass the summed confidence factor threshold is sent to the compression processor without an image snippet.

The Compression Engine employs a wavelet-based algorithm [14] to compress each image snippet, and formulates output contact messages in either of two types for output to a Queue Handler: a "contact" message that contains only contact data without snippet information and a "snippet" message that contains contact data with appended snippet information. The Queue Handler forms a prioritized output queue for each of these message types. The contact message queue is constructed to hold up to 200 contact-only messages, while the snippet message queue holds up to 20 snippet-image contact messages. The priority of the messages within each queue is based on an average of the CAD/CAC confidence factors for each contact cluster.

After either queue reaches its capacity, the lowest priority contact message within that queue is discarded and replaced by any subsequent higher priority incoming message having a higher average confidence factor. Periodic requests are received for data to be acoustically transmitted from the AUV to an instrumented buoy, which then relays the data via an RF transmission to a command operator on board a remote high-speed vehicle (HSV). When these requests are received, the Queue Handler sends out the highest priority contact within the appropriate queue. If the queue is empty when a request for data comes in, the request is ignored. Each time either queue is updated based on incoming messages or the removal of data from a queue, the contents of each queue are written to a file. The purpose of this queue file is to maintain the contents of the queues for re-initialization in the event that the data in the queues is lost due to power failure.

III. Description of CAD/CAC Fusion Experiment

Real-time CAD/CAC Fusion sea test demonstrations were performed during 2002 in a littoral coastal environment. The mine hunting operations were coordinated and conducted from the U.S. Navy platform HSV Joint Venture utilizing two REMUS AUVs, which operated within simulated minefields populated with inert MLOs. The AUVs were remotely operated by a team of operators from the Woods Hole Oceanographic Institute (WHOI) on board the HSV in conjunction with several navy divers. Each of the AUVs was outfitted with a different suite of sensors to accomplish a specific set of tasks in support of the mission objectives.

Two small boats were initially deployed to release two instrumented buoys at boundaries along the search area. Each buoy consists of an acoustic transducer on the in-water side and an RF transceiver on the in-air side. The primary function of these buoys is to provide a communications link between the operators aboard the HSV joint venture and the two REMUS vehicles during mine hunting operations. The secondary functions of the buoys are to provide the REMUS vehicles with a boundary of the survey area and the capability to localize themselves in a fixed coordinate system. Such a localization capability allows the REMUS vehicle to tag any MLOs with a fixed coordinate and also to reacquire a previously detected contact based on provided coordinates.

REMUS #1 was outfitted with port and starboard MSTL high frequency side scan sonar arrays to search the survey area and conduct real-time CAD/CAC operations, detecting mine-like objects located on the ocean floor in accordance with the operating concept illustrated in Figure 3.

The real-time CAD/CAC Fusion processing is programmed in an embedded 300 MHz Pentium processor installed in the AUV. After initial deployment, REMUS #1 obtained an initial coordinate fix by utilizing the previously deployed acoustic buoys and then began a search pattern consisting of preprogrammed tracks. While conducting

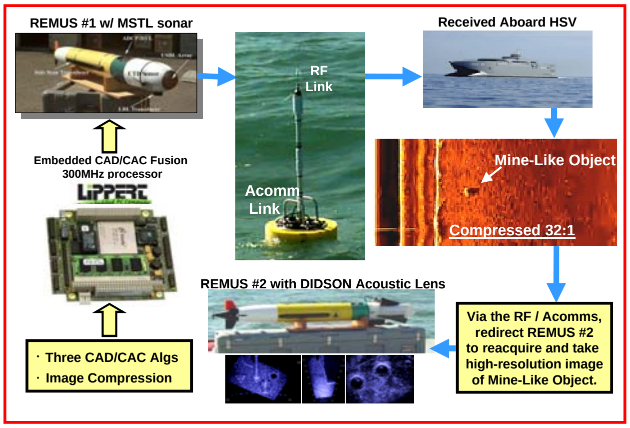

Fig. 3. Real-Time CAD/CAC Fusion Operating Concept.

these operations, this REMUS acquired port and starboard side-scan images, navigation data, fathometer records, and vehicle health and status. Each side-scan image was processed as discussed in Section II to generate the MLO contacts in real-time for report back to the command ship as the AUV continues its search.

REMUS #2 was outfitted with a Dual Frequency Identification Sonar (DIDSON) acoustic lens operating at very high frequencies [15,16] to further interrogate MLOs previously classified by the real-time CAD/CAC on board REMUS #1, and to provide images at significantly higher resolution for mine identification. Upon arrival at the survey area, REMUS #2 went into a holding pattern exterior to the track of the REMUS #1 mapping vehicle. As operators on board the HSV Joint venture received transmissions from the REMUS #1 consisting of information pertaining to MLOs, they designated MLOs for REMUS #2 to reacquire and identify as mines or non-mines.

IV. RESULTS

The real-time CAD/CAC Fusion was demonstrated at an ONR-sponsored littoral exercise during 2002. The mission spanned a total of 72 sidescan sonar images (port and starboard), in which there were 4 target opportunities. The REMUS #1 vehicle performing the search was commanded to transmit the top three clustered contacts from the contact queue at the end of each image collection/processing cycle, resulting in a total of 216 clustered contacts collected at the command ship during the area search. The top 4 highest priority contacts (highest average confidence factor) in this set of clustered contacts that were transmitted in real-time to the ship were the 4 valid MLO test targets. The geographic locations of these targets within the search area are illustrated by the red stars in Figure 4, which shows the plotted locations of the top 216 contact clusters, color-coded by cluster type.

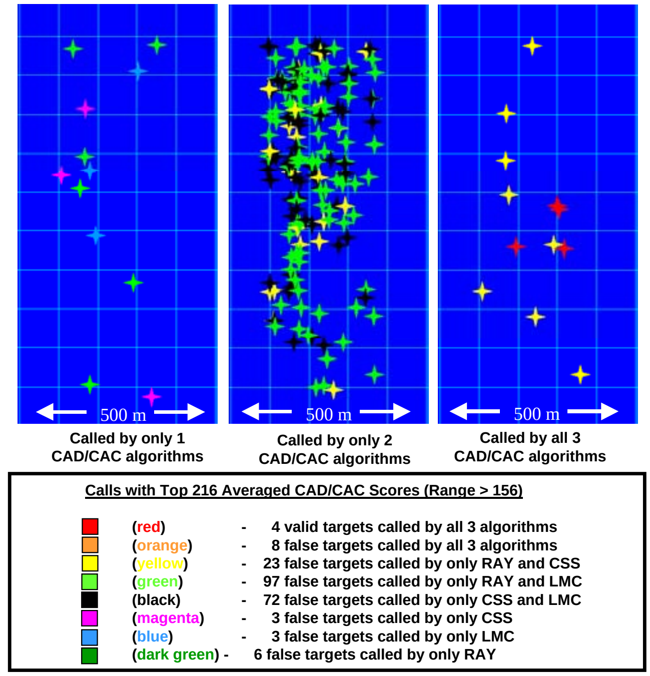

Fig. 4. Real-Time CAD/CAC Fusion Test Results.

The left side of the figure plots only those clustered contacts classified by a single CAD/CAC algorithm; the middle plot shows those contacts classified by 2 of the 3 CAD/CAC algorithms; and the right side shows the contacts classified by all 3 CAD/CAC algorithms. Post analysis of the recorded data resulted in the summary data of Table 1. The Raytheon, CSS, and Lockheed Martin CAD/CAC algorithms individually classified 788, 399, and 806 output contacts as MLOs over the entire mission, respectively. Each algorithm correctly classified all 4 targets. The individual algorithm classification outputs were clustered into a total of 1710 contacts, which includes the "singleton" contacts called by only one CAD/CAC algorithm. Since 4 of these were the targets, a total of 1706 false alarms were classified. An M-of-N fusion rule that requires at least 2 of 3 algorithms to call a contact reduces the total number of false alarms to 203, while still correctly classifying all 4 targets. This represents about a 2:1 reduction in false alarm rate relative to the best performing single CAD/CAC algorithm. As indicated by Figure 4, a 3-of-3 fusion rule maintains all 4 correct target classifications while further reducing the total false alarms to 8, a 50:1 reduction in false alarm rate relative to the best performing single algorithm.

Table 1. Summary CAD/CAC Fusion Test Results from Post Analysis.

CAD/CAC Fusion Algorithm	*Total # Contacts*	*# Targets*	*# False Alarms*	*FA Reduction*
Standalone CSS CAD/CAC	399	4	395	n/a
Standalone Raytheon CAD/CAC	788	4	784	n/a
Standalone LM CAD/CAC	806	4	802	n/a
2-of-3 Fusion	203	4	199	2:1
3-of-3 Fusion	12	4	8	50:1

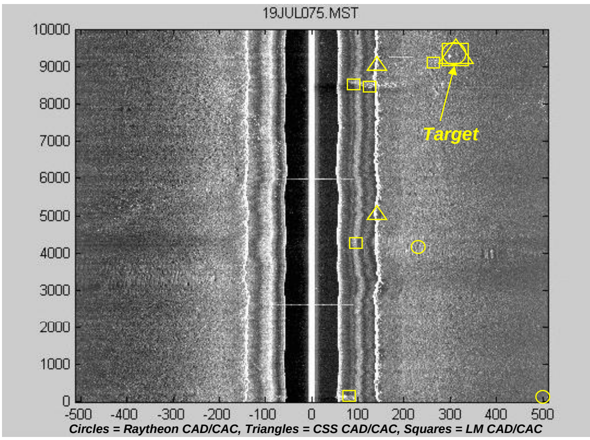

All 9 False Alarms Eliminated by Fusion While Retaining Target

Fig. 5. Example Sonar Image Showing Correct Classification of Mine-like Target.

Figure 5 shows an example sidescan sonar image from the test containing a valid target that all 3 CAD/CAC algorithms correctly classified, along with a combined total of 9 false alarms from the individual CAD/CAC algorithms. Note that all false alarms are singletons, which can be eliminated by the fusion processing through the use of a 2-of-3 fusion rule or a thresholding of the average confidence factor using a threshold level of 1/3. Figure 6 shows a second example test image containing 12 singleton false alarms, all of which are also readily eliminated through the fusion.

V. Discussion

The principal criteria for CAD/CAC performance are (1) maintaining a high probability of target detection and classification, and (2) minimizing the time duration between the time when sonar data is available for analysis and the time when information is compiled on potentially viable targets. The second criterion is related to the well-know criterion: minimizing the number of false target detections; i.e., the more false target detections there are, the more time it takes to do the analysis. However, if fast analysis software tools are available to allow an operator to quickly eliminate false detections made by CAD/CAC, then the number of false alarms can be somewhat high and not dramatically impact overall analysis time. In this section we would like to expound on two different ways to use real-time CAD/CAC to help reduce analysis time while maintaining a high probability of detection and classification.

First, real-time CAD/CAC results can be stored on-board the UUV and processed post mission. The moment the mission ends, CAD/CAC is complete and the operator can immediately review the CAD/CAC contacts. The obvious benefit of this use of real-time CAD/CAC is that the delays associated with performing the CAD/CAC processing post-mission are avoided. Moreover, since the operator need only review image snippets of potential targets vice the full imagery, a laborious and time-consuming process is eliminated. The process is further speeded by restricting operator review to only the image snippets of the detections with the top N fused scores, the review process is further speeded. N is typically selected to be proportional to the area covered. Note that the top-N scheme is loosely related to selecting a constant false alarm rate (CFAR) threshold, a strategy that is commonly used in Automatic Target Recognition. If the snippet review software tool is well designed, N can be quite large without impacting analysis time. (As an aside, with the future improvement and evolution of CAD/CAC, one will be able to select smaller values of N.)

Second, when real-time underwater communication is available, real-time CAD/CAC permits decision making to start with the first viable contact. One need not wait until the mission is completed. To maintain a low number of false alarms, real-time CAD/CAC thresholds can be increased at the cost of a slight reduction in probability of correct classification. The slight reduction in target classification rate is often offset by having the capability to know in real time that some targets are present; having such advanced knowledge can be advantageous. More complete analysis of all CAD/CAC contacts (using a lower threshold) can always be done post mission.

VI. Summary

The fusion of multiple CAD/CAC algorithms to reliably classify bottom mine-like objects in the littoral environment has been adapted to operate in real-time through the implementation of an effective hardware and software processing

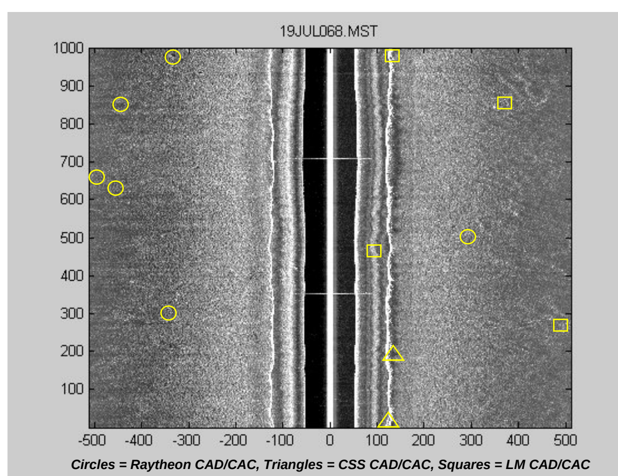

All 12 False Alarms Eliminated by Fusion

Fig. 6. Example Sonar Image Illustrating Elimination of all False Alarms Through Fusion.

architecture in the REMUS unmanned underwater vehicle. Sea testing conducted in July 2002 demonstrates that the system successfully employs the multiple CAD/CAC algorithms to correctly classify test mine-like targets while providing a significant reduction (up to 50:1) in false alarm rate relative to any single CAD/CAC algorithm. This system feature provides a valuable, timesaving automation aid to an operator who remotely receives the real-time classification outputs from the vehicle via acoustic/RF transmissions to/from a sonobuoy relay.

This real-time technology can be used to detect and classify objects in the cluttered littoral environment. Potential underwater applications include environmental monitoring, homeland security (waterway monitoring), underwater crime scene investigations, unexploded ordinance (UXO) detection, and archeological surveys.

Acknowledgements

The authors gratefully acknowledge the sponsorship of the Office of Naval Research (ONR 321OE, ONR 32MIW) for supporting this study. In addition, Tom Aridgides of Lockheed Martin is acknowledged for providing the outputs of his individual CAD/CAC classifier for use in the data fusion analysis. Roger Stokey, the REMUS software engineer at Woods Hole Oceanographic Institute, is also acknowledged for interfacing the REMUS vehicle to the CAD/CAC processor.

References

[1] Johnson, S. G. & Deaett, M. A., "The Application of Automated Recognition Techniques to Side-Scan Sonar Imagery," *IEEE J. Oceanic Engineering*, 19, No. 1, 138-144, January 1994.

[2] Ciany, C. M., Zurawski W., & Kerfoot, I., "Performance of Fusion Algorithms for Computer Aided Detection and Classification of Mines in Littoral Obtained from Testing in Navy Fleet Battle Exercise-Hotel 2000," Proceedings of SPIE'01, Vol. 4394, Orlando, Florida, 16-20 April 2001, pp. 1116-1122.

[3] Ciany, C. M. & Huang, J.,"Computer Aided Detection/Computer Aided Classification and Data Fusion Algorithms for Automated Detection and Classification of Underwater Mines," *IEEE Oceans 2000 MTS/IEEE Conference and Exhibition*, Vol. 1, Providence, R.I., 10 September 2000, pp. 277 –284.

[4] Ciany, C. M. & Zurawski W., "Application of Fusion Algorithms for Computer Aided Detection and Classification of Bottom Mines to Shallow Water Test Data," Proceedings of SPIE *AeroSense 2002, Detection and Remediation Technologies for Mines and Minelike Targets VI*, Orlando, Fl., 1-4 April 2002.

[5] G. Dobeck, "Algorithm Fusion for the Detection and Classification of Sea Mines in the Very Shallow Water Region using Side-Scan Sonar Imagery," Proceedings of SPIE'00, Vol. 4038, Orlando, Florida, 24-28 April 2000, pp. 348-361.

[6] G. Dobeck, "On the power of algorithm fusion," Proceedings of SPIE'01, Vol. 4394, Orlando, Florida, 16-20 April 2001.

[7] A. Aridgides, M. Fernandez, G. Dobeck, "Fusion of Sea Mine Detection and Classification Processing Strings for Sonar Imagery," Proceedings of SPIE'00, Vol. 4038, Orlando, Florida, 24-28 April 2000, pp. 391-401.

[8] A. Aridgides, M. Fernandez, G. Dobeck, "Side scan sonar imagery fusion for sea mine detection and classification in very shallow water," Proceedings of SPIE'01, Vol. 4394, Orlando, Florida, 16-20 April 2001.

[9] G. J. Dobeck, J. C. Hyland, L. Smedley, "Automated Detection/Classification of Sea Mines in Sonar Imagery," Proceedings of SPIE'97, Vol. 3079, pp. 90-110, Orlando, Florida, 20-25 April 1997.

[10] G. Dobeck, "Fusing Sonar Images for Mine Detection and Classification," Proceedings of SPIE'99, Vol. 3710, pp. 602-614, Orlando, Florida, 5-9 April 1999

[11] T. Aridgides, M. Fernandez, G. Dobeck, "Adaptive Clutter Suppression, Sea Mine Detection/Classification, and Fusion Processing String for Sonar Imagery," Proceedings of SPIE'99, Vol. 3710, pp. 626-637, Orlando, Florida, 5-9 April 1999

[12] L. Lam, C. Suen, "Application of Majority Voting to Pattern Recognition: An Analysis of Its Behavior and Performance," IEEE Trans on Systems, Man, and Cybernetics - Part A, Vol. 27, No. 5, September 1997, pp. 553-568.

[13] Ciany, C. M., Zurawski W., & Dobeck, G. "Application of Fusion Algorithms for Computer Aided Detection and Classification of Bottom Mines to Shallow Water Test Data From the Battle Space Preparation Autonomous Underwater Vehicle (BPAUV)," Proceedings of SPIE *AeroSense 2003, Detection and Remediation Technologies for Mines and Minelike Targets VIII*, Orlando, Fl., 21-25 April 2003.

[14] J. Walker & T. Nguyen, "Adaptive Scanning Methods for Wavelet Difference Reduction in Lossy Image Compression," Proceeding of ICIP, September 2000.

[15] E. Belcher, W. Fox, W. Hanot, "Dual Frequency Acoustic Camera: A Candidate for an Obstacle Avoidance, Gap-

Filler, and Identification Sensor for Untethered Underwater Vehicles," *IEEE Oceans 2002 MTS/IEEE Conference and Exhibition*, Vol. 4, Biloxi, MS., 29 October 2002, pp. 2124 –2128.

[16] E. Belcher, "Dual Frequency Identification Sonar (DIDSON)," APL/University of Washington web site: http://www.apl.washington.edu/programs/DIDSON/DIDSON.html