

MULTISPECTRAL AND HYPERSPECTRAL IMAGERY FOR OCEAN SEARCH AND RESCUE: DETECTION AND IDENTIFICATION OF BACKGROUND ANOMALIES

Dennis M. Silva, K. Mike Tao, Ronald Abileah, and Dennis E. Moellman
Applied Electromagnetics Laboratory
SRI International
Menlo Park, California 94025

ABSTRACT

Images collected by an airborne 48-channel imaging spectrometer are used to test the effectiveness of multispectral image processing techniques for detecting and identifying search and rescue (S&R) targets in a simulated ocean accident scene. The “push broom” sensor employs successive rasterlike scans of the ocean surface, spanning the visible portion of the spectrum from 430 to 830 nm. Statistical and physics-based approaches are employed to construct a background model that describes ocean surface clutter. A decluttering algorithm developed by SRI International, based on the background model, is used to remove the background clutter in the ocean imagery collected by the sensor. Pixels not exhibiting the spectral characteristics of illuminated ocean are intensity enhanced to appear as bright variations in the processed imagery, rendering the objects of interest to an S&R mission much easier to detect. In this context, we discuss the efficacy of multispectral imagery for clutter rejection, detection, and anomaly cueing. Our study indicates that automated image processing using data collected by multispectral or hyperspectral sensors could be a valuable tool for ocean S&R applications, allowing detection and identification to be achieved at higher altitudes and improving search rates when time is critical.

Key Words: multispectral, hyperspectral, search and rescue, ocean clutter, clutter removal

I. INTRODUCTION

The purpose of this work is to extend the application environment of a real-time ocean-surface decluttering algorithm previously developed by SRI International (SRI) [1,2]. The new application is search and rescue (S&R). The objective is to differentiate normal ocean clutter from S&R objects of interest at a single-pixel level. This capability would allow S&R operations to be conducted at significantly higher altitudes than is the current practice, which is often limited by the unaided abilities and endurance of human observers. The operational altitude of the S&R platform is significant in that it is directly related to search rate.

The S&R objects of interest in this study are assumed to occupy an area on the order of 1 m^2 (such as a person in the water). The sensor pixel resolution should be of the same order. For single-pixel detection, sensitivity loss is 6 dB per doubling of the pixel scale over the size of the object. In the context of spectral processing, smaller pixel sizes offer no advantage either since they only add to the processing without providing any gain in detection. Detection gains via spatial processing are potentially available from the exploitation of higher spatial resolution data, but at a significant computational cost. As a consequence, 1 m spatial resolution is the appropriate scale for S&R applications. Currently, multi- and hyperspectral imaging systems are capable of 1 m spatial resolution at altitudes of at least 2 km. HYDICE is an example of one such instrument. Objects of larger size will require lower spatial resolution, allowing for enhanced search rates since the search aircraft will be able to fly at even higher altitudes.

The requirement of 1 m spatial resolution impacts other operational parameters in S&R applications. For current push-broom-type sensors, the cross-track field of view is on the order of 250 to 500 m. The scan rate determines the aircraft speed; therefore, typical scan rates of 50 to 100 Hz limit aircraft to speeds of 100 to 200 knots. Thus, current sensors can evaluate $\sim 15,000$ to $30,000$ pixels per second, or, equivalently, 50 to 100 km^2 (20 to 40 mi^2) per hour.

By contrast, the *National Search and Rescue Manual* (www.uscg.mil, search query “MSN”) limits the search rate for an individual in the water by fix-wing aircraft to about 50 km^2 (20 mi^2) per hour under ideal conditions. For the most part, the search rate is determined by the swath width, with swath widths increasing as S&R target sizes increase. As wind and sea conditions worsen, however, the search rate must be reduced significantly to accommodate several operational requirements imposed on the swath width. For example, if surface wind speeds exceed 15 knots, the swath width must be reduced by a factor of 2; they are reduced by an additional factor of 2 if surface wind speeds exceed 25 knots. The swath width must also be adjusted for visibility, aircraft speed, and crew fatigue. We anticipate that with the use of multispectral processing techniques to remove extraneous interference from background light, most of the swath-width reductions can be reclaimed and the search rate can be restored without affecting the detectability of S&R targets.

In general, detection will be limited by the ability of the processing to remove background light. The ambient light in an ocean scene consists of four components: (1) upwelling light produced by backscattered solar light interacting with subsurface water molecules; (2) sky background—skylight reflected from the surface and scattered from below; (3) glint—specular reflection of sunlight off the surface; and (4) whitecaps—surface foam created by breaking waves. The SRI ocean decluttering algorithm decomposes the light intensity in each pixel into these components and subtracts them, leaving only residual noise from the sensor and reflected light that is spectrally different from the natural background (i.e., S&R target candidates).

The SRI algorithm was previously developed for detection of submerged objects [1,2]. This study extends the approach to floating objects that would be considered of interest in S&R surveillance. Section II covers results obtained with the SRI algorithm on hyperspectral imagery collected over a simulated ocean crash site. In Section III, we discuss the spectral content of the imagery. In Section IV, the results are analyzed in the context of a classification problem. Concluding remarks are provided in Section V.

II. RESULTS OF IMAGE DECLUTTERING AND TARGET CUEING

To provide a preliminary look at the effectiveness of spectral-based optical processing for S&R applications, a proof-of-concept test was carried out. The experiment was a joint effort conducted by the U.S. Coast Guard, the Space and Naval Warfare Systems Center (SPAWAR, San Diego, California), and SRI (Washington, DC, and Menlo Park, California). The exercise was conducted off the coast of San Diego over a two-day period. The Coast Guard provided both the services of a patrol boat and debris simulations of the type one might expect from a mishap at sea: a pilot and aircraft wreckage and other equipment, all floating on or near the surface. This collection of objects (debris array) was tethered together and attached to the patrol boat for towing. The simulated pilot, along with the other elements of the debris array, were predominantly yellow, red, and orange. Figure 1 is a photograph of a crewman deploying one of the debris arrays used in the exercise. The services of the multispectral sensor and platform were arranged by SPAWAR, San Diego. The sensor that performed the collection captured 48-channel spectra from 430 to 830 nm in even bandwidth steps of 8.33 nm per channel and had a spatial resolution $\sim 1 \text{ m}^2$. The sensor platform overflew the deployed debris arrays several times on both days of the experiment. SRI was responsible for the analysis and results described in this paper.

The SRI decluttering process was applied to a set of images recorded over the simulated S&R scene. Three ocean scenes were used in this study. Two of the images were recorded on the first day of the two-day exercise, and the third image (Figure 2) was recorded on the following day. Similar results were obtained for all three images, so, for brevity, only the third image is discussed in detail.

The bright feature in Figure 2 is the U.S. Coast Guard patrol boat that was supporting the experiment. The patrol boat appears somewhat distorted because the along-track to cross-track pixel ratio was not one-to-one. On the day of the collection, the weather was mild and the sea was relatively calm. The whitecap population would be expected to be well below 1% of surface coverage [3–5]. The original image shown in Figure 2 contains the aforementioned patrol boat, a simulated debris field, and a small boat. The result obtained after decluttering is shown in Figure 3, and Figure 4 provides a thresholded version of Figure 3. The three images used in this study shared the same

(operational) threshold value that was set just above the largest background value found in the processed imagery. The patrol boat, small boat, and debris field are called out in the processed data. As can be observed in Figure 3, the image appeared to be made up of two halves. The artifact was caused by the use of dual sensors, with each half of the image recorded by a different sensor. The left-right difference in the combined image was caused by a difference in the noise characteristics associated with the two sensor components. For this work, the data from the two sensors were processed together as a single scene.

The original image shown in Figure 2 is dominated by ocean-surface glint, which the decluttering algorithm removes very effectively. The wake left by the small boat is another interesting feature in the image. It has the same spectral characteristics as foam normally associated with breaking-wave whitecaps. Although bright in the original image, it is greatly reduced in the processed image because the decluttering process effectively eliminates whitecaplike features.

Having removed or reduced the background clutter, the most significant pixels are those associated with the patrol boat, the debris field, the small boat, and a pixel-sized spot located to the right of and below the small boat. Pixels exceeding the operational threshold (Figure 4) are to be interpreted as nonbackground anomalies.

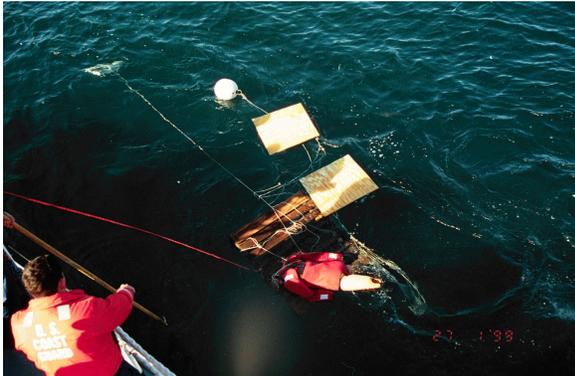


Figure 1. Sample of debris array used as targets for the search and rescue study.

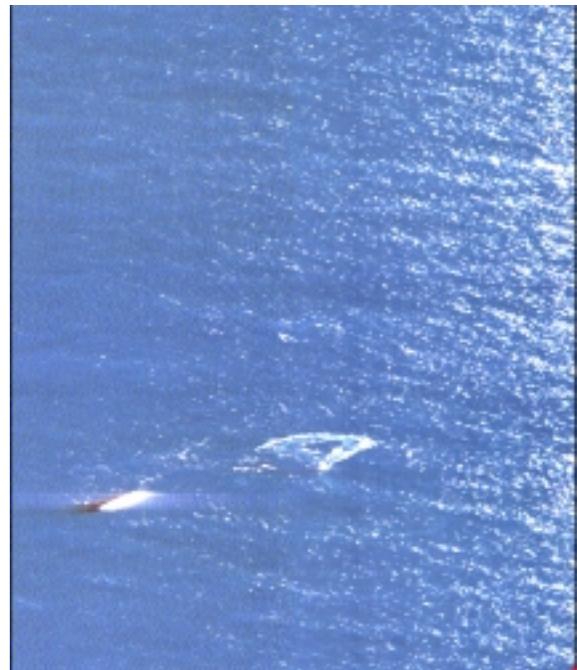


Figure 2. Search and rescue scene used for analysis.

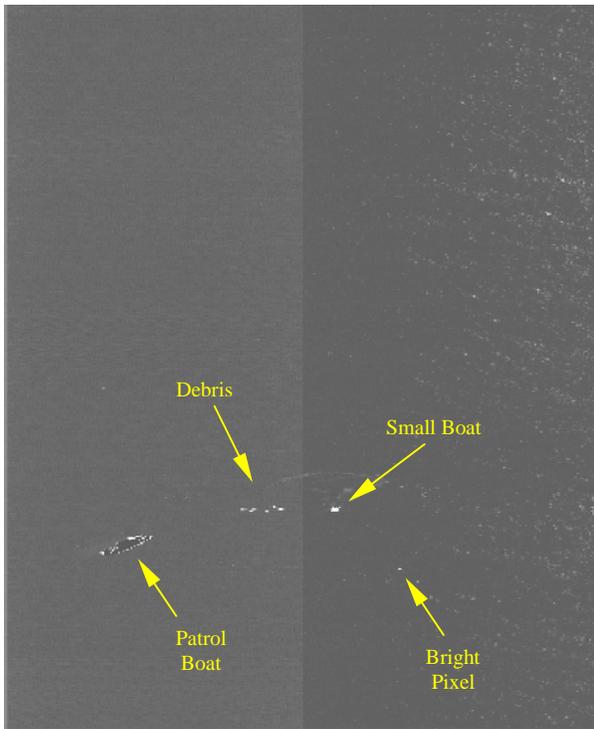


Figure 3. Bright pixels remaining after processing of search and rescue image.

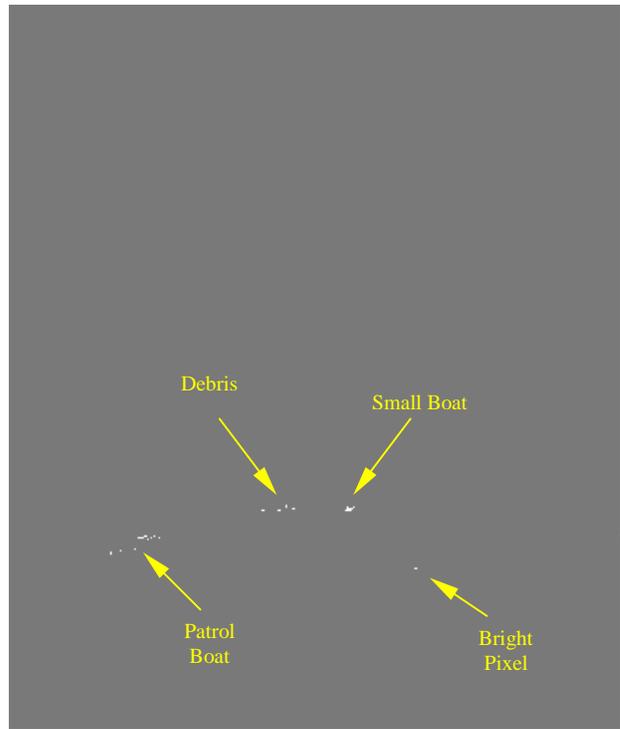


Figure 4. Search and rescue image after thresholding the processed image.

III. COMMENTS ON THE SPECTRAL CONTENT OF THE OCEAN IMAGERY

The pixels that appear dark in the processed image are those that are better explained by the background clutter model. Figure 5a shows a background spectrum from a typical sun-glint-free pixel. As is to be expected, the spectrum is peaked in the blue-green. Glint spectra are similar but contain more red; an example of one such spectrum is shown in Figure 5b. The typical spectra associated with whitecaps are, as we would expect, white or nearly flat [6] over most of the range of visible light (Figure 5c).

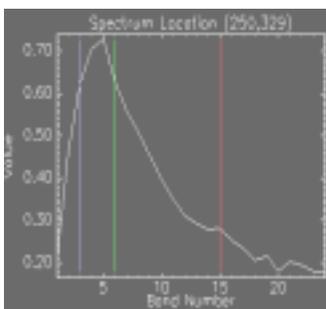


Figure 5a. Example of upwelling-light spectrum.

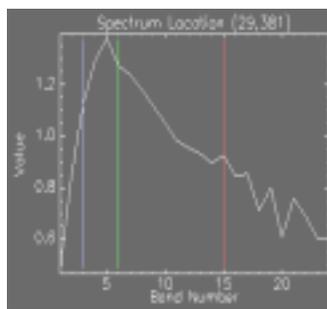


Figure 5b. Example of glint spectrum.

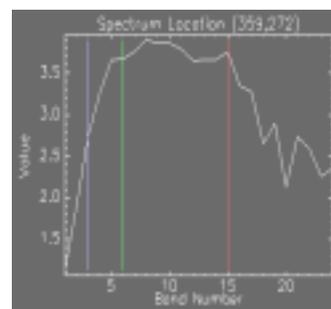


Figure 5c. Example of whitecap spectrum.

The pixels called out in the scene after processing have spectral properties that are very different from background spectra. Some examples of single-pixel spectra associated with four of the objects in the debris field are provided in the blowups in Figure 6. These spectra appear to be dominated by orange and red spectral components that have caused them to be flagged as nonbackground anomalies and, when compared with the brightness scale of upwelling light (Figure 5a), are very intense.

Samples of single-pixel spectra associated with the small boat and the aberrant pixel are shown in Figure 7. The pixels associated with the small boat assumed more than one characteristic shape, but the most consistent spectral shape was the orange-red coloration of the boat itself. The isolated pixel appeared orange, and had a spectrum nearly identical to that of what was presumed to have been an international orange float observed on the previous day within the deployed debris field. None of the pixels called out by the processing in this scene would be mistaken for normal ocean background.

IV. EXTENDED ANALYSIS OF RESIDUALS AND SPECTRUM CLASSIFICATION

In Section II, an operational threshold was used to flag the pixels of interest. In this analysis, we studied the pixel population at lower threshold values. The purpose of such an exercise is to determine the spectral characteristics of statistical outliers that were relatively bright but did not exceed the operational threshold. It serves to define the extent to which the background model fits the background data and to provide a basis for spectrum classification.

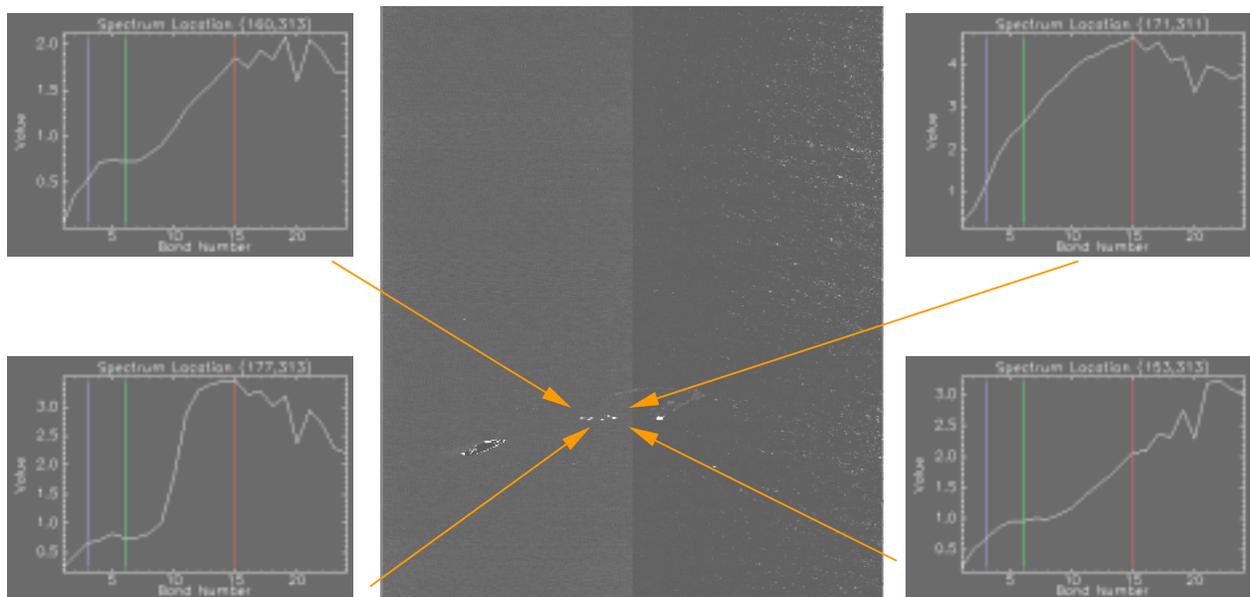


Figure 6. Assorted spectra associated with the objects of interest in the debris field.

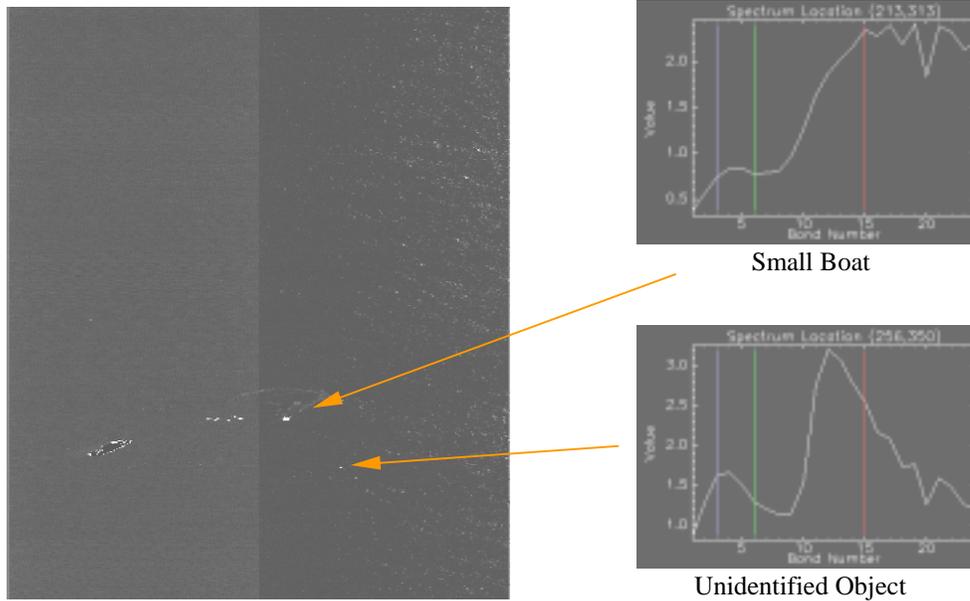


Figure 7. Spectral sample of pixels associated with the small boat and the unidentified object.

The analysis was performed using two of the S&R images. In this exercise, the original 48-channel data were binned into 24 channels by pair-wise summing adjacent channels. These 24 channels represented the maximum spectral resolution available because of the procedure used to calibrate the data. The decluttering algorithm was applied to the images that included the image shown in Figure 2 and an image recorded on the previous day (Figure 8). The 24-channel decluttered output obtained from the image shown in Figure 8 was first scanned for statistical outliers at a threshold significantly lower than the operational threshold used for cueing (Section II). Since the decluttered data had the appropriate bell-shaped distributions, standard deviations could be estimated from the pixel populations following decluttering to provide a reasonable basis for outlier determination. Pixel intensity deviations that significantly exceeded the standard deviation were declared outliers and plotted in a lower (three-) dimensional feature space. The three axes of the feature space represented the pixel intensities of channels that corresponded to wavelengths in the red, green, and blue portions of the visible spectrum, denoted in Figures 9a and 9b by the axis labels R, G, and B, respectively.

Using the image shown in Figure 8, the pixel outliers after processing were first identified. With the patrol boat and its immediate surroundings omitted from consideration, the aberrant pixels appeared to fall into two classes. One class was associated with an apparent off-color surface film that can be seen at the top of the image. The second class was associated with the residual glint regions that were not bright enough to survive the operational threshold imposed in Section II, but were bright enough to pass the relatively low threshold criterion of this analysis. Inasmuch as a surface film would have to be regarded as interesting in an S&R scenario, there was only one anomaly class that could be associated with the natural background. Given the unlikely presence of whitecaps under the mild wind and sea conditions at the time of the collection and the spectral properties of the residuals, it is likely that the background anomaly class was associated with residual glint that was not completely removed by the decluttering algorithm. Although not completely removed, the decluttering algorithm performed well enough to render the residual glint unimportant in the processed image. When viewed in RGB feature space (Figures 9a and 9b), it is evident that the power to discriminate between these two classes could be provided by the individual pixel spectra as well as by their principal directions, along which the spectra of the two anomaly classes tended to vary most in this image.

Proceeding with a spectral processing approach based on self training, the two background anomaly classes derived from the image shown in Figure 8 could be considered library elements to be used to identify features in other imagery. The image shown in Figure 2 was used for such a test. The images (Figures 2 and 8) were recorded on different days, under different environmental conditions. Using a nearest-neighbor (NN) classifier [7] with an extended capability to classify spectral patterns as unknown, the pixels in the Figure 2 image were classified using the library constructed from Figure 8. The outlier pixels were classified into one of three classes: surface film, glint, or unknown. Since the spectra within each class appeared to be described adequately by a Gaussian model, the metric used to determine the “distance” between a given spectral pattern and a known class was given by

$$(x - m_i)^T Q_i^{-1} (x - m_i) ,$$

where m_i and Q_i were the mean feature vector and the covariance matrix of a particular anomaly class i , and x was the input pattern vector being classified.

The metric represented by the equation above is sometimes referred to as Hotelling’s transform, the square root of which corresponds to the standard deviation measure (in units of sigma) typically associated with single-variable statistics. If a 2-sigma threshold is chosen, the classifier will classify the input pattern according to its NN class if it is within a 2-sigma distance from that class; otherwise, it will declare the input pattern a novelty (i.e., unknown) and worthy of further investigation. (The NN classifier is considered as a viable approach for its simplicity and robustness. With the prior probability of each class reliably known, the NN classifier can be readily extended to become a Bayesian classifier.)



Figure 8. Image used for analysis of outliers.

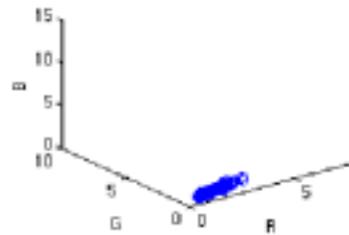


Figure 9a. Distribution of outlier pixels associated with surface film in RGB feature space.

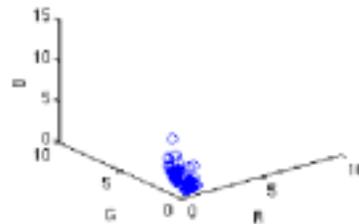


Figure 9b. Distribution of outlier pixels associated with glint in RGB feature space.

As can be seen in Figures 3 and 4, there are a number of man-made objects in the scene. Applying the classifier built around the two-class common background library derived from the image in Figure 8, the man-made objects in the scene are called out as falling into the unknown class at the 2-sigma threshold setting.

When applied to imagery taken under significantly different conditions, this classification procedure can be made even more robust by normalizing the pixel spectral vector and by using the principal directions of the spectral variations. The latter features appear to be insensitive to the varying imaging conditions encountered in this study.

V. CONCLUSIONS

The SRI ocean-surface decluttering algorithm was applied to a number of ocean-imaging scenes. The emphasis of the study was directed at an assessment of the performance of the algorithm on data intended to replicate the imagery anticipated in a typical S&R scenario. The algorithm was able to cue the nonbackground target-object pixels uniquely in all three S&R images as anomalies. In the sample image provided (Figure 2), there was only one aberrant pixel that was located at a significant distance from the debris field. This pixel could not be classified as normal ocean background, having a spectral signature characteristic of one of the international orange floats used in the experiment. For all three images analyzed in this study, the only other source of cueing (outside the test debris areas) was caused by an apparent surface film or dye marker. Cues caused by surface films, however, qualify as nonbackground anomalies. The spectral character of this anomaly was unique and easily identified as the same as the surface film observed in much greater abundance in another part of the image.

In conclusion, we estimate that our multispectral cueing method can increase search rate by a factor of 2 or more, depending on sea conditions, operator fatigue, and so on. Additional research is required, however, since the only precise way to gauge the improvement in search rate is with testing. The technology is mature and ready for such testing.

ACKNOWLEDGMENTS

The authors would like to thank the U.S. Coast Guard for their assistance in setting up the S&R experiment. The authors would also like to thank Dr. Jon Schoonmaker of SPAWAR, San Diego, for arranging for the collection and dissemination of the data used in this study, and Mr. Robert Chiralo of SRI International for his informative comments.

REFERENCES

- [1] D. M. Silva and R. Abileah, "Two Algorithms for Removing Ocean Surface Clutter in Multispectral and Hyperspectral Images," *Ocean Optics IV*, Kona, Hawaii, 1998.
- [2] D. M. Silva and I. E. Abdou, "Removing Ocean Surface Clutter in Multispectral and Hyperspectral Imagery," *Proc. SPIE*, **3717**, pp. 135–147, 1999.
- [3] E. C. Monahan, "Oceanic Whitecaps," *J. Phys. Ocean.*, **1**, pp. 139–144, 1971.
- [4] J. Wu, "Oceanic Whitecaps and Sea State," *J. Phys. Ocean.*, **9**, pp. 1064–1068, 1979.
- [5] I. I. Strizhkhin, "Features of the Distribution of Foam on the Sea Surface," *Meteorologiya i Gidrologiya*, **7**, pp. 92–97, 1988.

- [6] P. Koepke, "Effective Reflectance of Oceanic Whitecaps," *Applied Optics*, **23**, No. 11, pp. 1816–1824, 1984.
- [7] T. M. Cover and P. E. Hart, "Nearest Neighbor Pattern Classification," *IEEE Transactions on Information Theory*, **IT-13**, No. 1, pp. 21–27, 1967.