UCRL-JC-114623
PREPRINT

Data Fusion for the Detection of Buried Land Mines

Gregory A. Clark, Sailes K. Sengupta, Paul C. Schaich, Robert J. Sherwood,
Michael R. Buhl, Jose E. Hernandez, Ronald J. Kane, Marvin J. Barth,
David J. Fields, and Michael R. Carter

This paper was prepared for submittal to the
Substance Identification Technologies
Innsbruck, Austria
October 4-8, 1993

October 1993

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Data Fusion For The Detection Of Buried Land Mines

Gregory A. Clark, Sailes K. Sengupta, Paul C. Schaich, Robert J. Sherwood,
Michael R. Buhl, Jose E. Hernandez, Ronald J. Kane, Marvin J. Barth,
David J. Fields, and Michael R. Carter

Lawrence Livermore National Laboratory
7000 East Avenue, L-156, Livermore, CA 94550

Please send correspondence to:
Dr. Gregory A. Clark
Lawrence Livermore National Laboratory
7000 East Avenue, L-156, Livermore, CA 94550
Telephone: (510) 423-9759
FAX: (510) 422-3013
E-Mail: clark9@llnl.gov

Abstract

We have conducted experiments to demonstrate the enhanced detectability of buried land mines using sensor fusion techniques. Multiple sensors, including visible imagery, infrared imagery, and ground penetrating radar, have been used to acquire data on a number of buried mines and mine surrogates. We present this data along with a discussion of our application of sensor fusion techniques for this particular detection problem. We describe our data fusion architecture and discuss some relevant results of these classification methods.

1. Introduction

The goal of this work is the detection and location of buried objects, specifically land mines, given images obtained from a suite of remote sensors. Past research has shown that it is extremely difficult to distinguish buried mines from background clutter in images obtained from any single sensor. It is hoped that information fused from a suite of different sensors will provide better detection reliability, because the variety of physical properties separates more clearly targets from clutter and background.

We are acquiring data over the relevant breadth of the electromagnetic spectrum, including two infrared sensors (5 micron and 10 micron wavelengths), a ground penetrating radar (GPR) of the short pulse, wideband synthetic aperture type, UV, visible and near IR images. The visible and near IR data is from a high resolution imaging spectrometer. The detection system uses advanced algorithms from the areas of automatic target recognition (ATR), computer vision, signal and image processing, and information fusion. The system allows the use of both physical principles and image processing for image interpretation.

This work is application research in progress. The individual algorithms used are advanced, but mostly known, and the novelty of the work lies in the combination of the algorithms and their application to the very difficult and important problem of detecting buried land mines. To date, no successful operational system exists for airborne standoff detection of buried mines. We believe that data fusion applied to simultaneous observations with a variety of different sensors may provide a reliable signature for buried mines. This paper focuses on the fusion of images from two infrared sensors with different band passes.

Section two describes the data which was used for analysis. Section three describes the data fusion and the automatic target recognition. Section four describes the results obtained in a clutter free environment. Section five describes the results obtained in a cluttered environment. Lastly, sections six through eight present conclusions, acknowledgments, and references, respectively.

2. Experiments and Measurements

The data set described here was taken at two separate sites. The Lawrence Livermore National Laboratory (LLNL) site was a clutter free site in which mine surrogates were buried in sand and clay plots with no covering vegetation. There were a total of 18 mine surrogates of different types and 9 holes with no mine surro-

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

gate. The Nevada Test Site (NTS) is located in the Nevada desert. Mines were buried with no deliberate alteration of the native vegetation. There are about 500 real mines and mine surrogates, both plastic and metal, buried at this site at depths between 1 and 10 cm. In both locations, care was taken to avoid leaving obvious surface signatures. At the NTS site it is not possible to locate mines visually. The data from the LLNL site was acquired in both the daytime and the nighttime. For the NTS site, data was only acquired in the nighttime.

Thermal infrared data was taken with two Agema 880 scanners covering the 3-5 μ m (S) and 8-12 μ m (L) IR regions. The instruments are capable of resolving temperature differences equivalent to 1°C. The images were taken from a tower or lift truck at heights from 13 m to 20 m. The 140 by 280 pixel field of view of the cameras covered an area 5m - 7m square, with ground resolutions between 2-6 cm. The data analysis discussed below will be primarily based on the fusion of these two bands of infrared imagery.

The ground penetrating radar (GPR) data[1] was taken at both sites using pulsed ultra-wide band systems. The data was processed as SAR imagery to provide images which could be coregistered with other data. While the radar systems do show signatures of buried mines, the analysis of data sets which include the GPR is still in progress at the time of this writing.

In addition to panchromatic visible data, UV and hyperspectral visible/near IR data[2] is also being acquired. The UV data is taken with a intensified camera with a UV sensitive photocathode and appropriate band pass filters to cover the 310-330 nm region. The hyperspectral data is taken with a line imaged grating instrument using a CCD as the imaging device. Its coverage is variable over the visible and near IR from 425 to 850 nm. The channel bandwidth is 3.4 nm in each of 125 channels. This data will be acquired in October 1993.

3. Data Fusion and Automatic Target Recognition

We use a supervised learning pattern recognition approach to detecting the metal and plastic land mines buried in soil. The overall process consists of four main parts: preprocessing, feature extraction, feature selection, and classification. These parts are used in a two step process to classify a sub-image. The first step, referred to as feature selection, determines the features of sub-images which result in the greatest separability among the classes. The second step, image labeling and post-processing, uses the selected features and the decisions from a pattern classifier to label the regions in the image which are likely to correspond to buried mines. We process images using a SUN Sparc 2 and the VISION software package written at LLNL (the primary author is J. E. Hernandez). VISION is an object-oriented package, and it runs under Franz Allegro CL, which implements the Common Lisp Object System (CLOS) [12].

3.1. Image Preprocessing

The preprocessing of images consists of operations on the image prior to feature calculation. This includes registration of images so that corresponding spatial position occur in corresponding image pixels, and calculation of pseudo-images from existing images in order to isolate specific physical processes into specific images.

Registration: The multiple images of the scenes do not, in general, superimpose correctly, due to fundamental sensor differences (scalings, fields of view, etc.) and sensor geometric distortion (barrel distortion, etc.). We manually identify fiducial markers on the ground to be used as control points for a perspective warping algorithm, which performs translation, rotation, scaling and perspective corrections to the images[3,4]. We obtain an ensemble of corrected images which can then be processed pixel-wise with the assurance that pixels in the various images correspond very closely with same point on the surface.

A separate geometric mapping operation is performed on each image. For each image to be mapped, we compute two mapping polynomials (one to map the X coordinates of the control points to their target locations, and the other to map the Y coordinates of the control points to their target locations). Because our images require perspective corrections, we used a perspective mapping algorithm [24-27]. This algorithm performs translation, rotation, scaling and perspective correction. We used four control points and their target locations in each image to determine the mapping. Once the mapping polynomials have been computed, each image can be geometrically mapped using its mapping polynomials. The mapping is performed by iterating through the pixels in the corrected image. At each pixel location in the corrected image, we compute the pixel coordinates of the corresponding pixel in the uncorrected image using the inverse mapping functions. In general,

these computed coordinates fall between pixels in the uncorrected image, so we use bilinear interpolation among the four nearest neighbors in the uncorrected image to compute the pixel value. A pixel of that value is inserted at that location in the corrected image. We then step to the next pixel in the corrected image and repeat the process. The result is the geometrically corrected image.

Calculation of Pseudo-images: If certain physical parameters which are more likely to contain signatures of targets can be emphasized in an image by combining data from more than one sensor, then this pseudo-image can be calculated and treated as an independent image. While sophisticated classification algorithms may infer these relationships from the original images, they may require a larger number of samples than are available to do so accurately. Hence, we have calculated two such images, called "emissivity" (e) and "temperature" (T). Emissivity is defined as the natural logarithm of " S/L ", and temperature is defined as the natural logarithm of " $L^2/S \cdot S_{\text{mean}}/L_{\text{mean}}^2$ ", where L is the long wavelength band, S is the short wavelength band, S_{mean} is the mean of the short wavelength band, and L_{mean}^2 is the mean of the squared long wavelength band. Emissivity should be insensitive to difference of temperature of a single substance, but should show contrast between different substances. Temperature should isolate differences in temperature of a substance and be less sensitive to differences of emissivity.

3.2. Feature Extraction

After the pre-processing has been performed, the sub-image samples in the image are selected. These sub-image are square regions ($N \times N$) which have dimension (N) of approximately the diameter of a mine. This allows the majority of the pixels to be contained within the mines. This value differed for the experiments presented in sections four and five. This was due to differences in pixel resolution in both cases.

Given these sub-images, we compute a vector of statistical features from the pixel values in the sub-images. For the results in this paper, we use amplitude histogram features and spatially-dependent features, including texture features [4, 5].

3.3. Feature Selection

Human experts generally classify objects based on a very few of the most important attributes in the image. The fundamental function of the feature selection process is to select the most useful information from the representation vector and present it in the form of a relatively low-dimensional pattern vector removing any redundant and irrelevant information which may have a detrimental effect on the performance of the classifier. A useful by-product in the process is knowledge about the discriminatory potential of the features and the associated highest achievable performance for a given set of features. Statistical decision theory tells us that the probability of misclassification is a decreasing function of the number of features provided, if the sample size is very large. In practice however, only a small number of training sets is available and estimation errors are no longer negligible. Since the number of parameters and the associated estimation errors increase rapidly with dimension, it may be advantageous to sacrifice some useful information in order to keep the number of these parameters to a minimum.

An important goal in our work is to use feature selection techniques to choose the subset of features that contribute most to correct classification. In order to accomplish this, features must be extracted from known sub-image samples, providing a known database. We gain two main benefits from feature selection. First, we wish to minimize the computational complexity of our processing algorithms, so they can eventually be implemented in "real time." Second, we wish to determine which sensors are the most important for classification. By rank ordering the features according to their importance for classification, we are able to eliminate from consideration sensors which do not contribute significantly. Feature selection is typically accomplished by computing a distance measure which is the sum of probabilistic distances between all pair-wise combinations of classes [5,6]. Commonly used algorithms include Branch and Bound, Sequential Forward Selection, and Sequential Backward Selection [5,6].

3.4. Classification

The pattern recognition problem is difficult because various sources of noise distort the patterns, and often there exists substantial variability among the patterns belonging to the same class [13-22]. "Rules" based on the premise such as "buried targets are warmer than the background during night" can be constructed providing a data-driven threshold for the purpose of classification. Our initial work focused successfully on this method [9-11, 23]. However, rules of this type are in effect modified parallelepiped classifiers that do not take into account the covariance structure of the classes. The rule-based approach is valuable, particularly as pre-processing step, but we have found that the supervised learning approach is generally more robust to varying data scenarios.

In our studies, we have used a variety of classifiers, including the nearest neighbor classifier [6], and the probabilistic neural network [8, 16].

4. Supervised Learning Results in a Clutter-Free Environment

The data which was processed for these results was acquired at the LLNL site. This site contained both clay and sand plots. Additionally, the data was acquired in both daytime and nighttime.

In this experiment, we defined a two-class problem. The first class (called "Mines") corresponds to buried mines and filled-in holes. Filled-in holes were combined with the buried mines because some preliminary results indicated that it is very difficult to differentiate them using thermal IR. Moreover, it is safe to assume, in a typical operational situation, that wherever a filled-in hole exists, a mine may also exist. The second class (called "No Mines") corresponds to background, clutter, and fiducial markers. Although these plots visually appear clutter free, there was some thermal clutter. This class was defined as regions in the image which were in the background but had thermal intensity values similar to mines.

The sub-images which were sampled from these images were of size 21. This size was selected because it was approximately the diameter of a mine. For the results discussed below, only the amplitude features were generated. This features consisted of mean, standard deviation, skewness, kurtosis, energy and entropy.

Because daytime and nighttime images have differing thermal characteristics, they were evaluated separately. Additionally, the clay and the sand images had differing thermal characteristics, so they were also evaluated separately. In an attempt to determine whether features which represent temperature and emissivity characteristics resulted in better detection probabilities than features which simply represent long and short wavelength intensity values, the four case given above were repeated once for long and short wavelength features and once for temperature and emissivity features. Therefore, eight cases were studied: daytime, clay, short and long; nighttime, clay, long and short; daytime, sand, short and long; nighttime, sand, long and short; daytime, clay, temperature and emissivity; nighttime, clay, temperature and emissivity; daytime, sand, temperature and emissivity; and nighttime, sand, temperature and emissivity.

4.1. Feature Selection

Because the sample set was small, conventional feature selection techniques, such as the branch and bound and the sequential forward selection algorithms, could not be applied. Since the ultimate goal of this work was to detect mines, the probability of detection was used as a means of defining class separability for features subsets. This was computed by training a classifier, which in this case was the nearest neighbor classifier, with all the feature vectors except for one feature vector was held out for testing. Following this the unknown vector was then classified. This was repeated for all the training vectors, and probability of detection was computed. Since each feature vector could consist of up to twelve features (six histogram features • two images), there were 4095 ($2^{12} - 1$) possible features subsets tested for each case. The feature subset which provided the highest probability of detection were selected as the "optimal" subset. The image labeling along with the post-processing are described in the next section.

4.2. Image Labeling and Post-Processing

After the feature selection had been performed, the "optimal" subset of features were computed for every pixel in the testing images. These feature vectors were then classified as either a "Mine" or a "No Mine" using a nearest neighbor classifier. In an attempt to "clean up" this binary image, a morphological "opening" operation was performed. Following this, a region growing algorithm was applied to isolate "Mine" regions. Regions which were very small (less than 1/4 the area of a mine) were eliminated from further consideration. The centroids of the remaining regions were estimated mine locations. If the estimated location was within a mine radius of the actual location, the detection was considered a success. Otherwise, the detection was considered a false alarm. The feature selection results and the detection results for all eight cases are presented next.

4.3. Results

Results of using long and short wavelength features are given in Table 4.1, and results of using temperature and emissivity features are given in Table 4.2. The "probability of detection" is calculated as the ratio of the number of actual mine locations detected to the total number of actual mines. The "probability of false alarm" is calculated as the ratio of the estimated mine locations which do not lie within a mine radius of an actual mine to the total number of estimated mine locations. These ratios are given below the probabilities in Tables 4.1 and 4.2.

Time of Day, Terrain, Image Pair	"Optimal" Feature Subset	Probability of Detection	Probability of False Alarm
Daytime, Clay, Long & Short	mean (long) standard deviation (long) skewness (long)	100.0 (18/18)	0.00 (0/18)
Nighttime, Clay, Long & Short	mean (short)	88.89 (16/18)	0.00 (0/16)
Daytime, Sand, Long & Short	standard deviation (long) kurtosis (long) standard deviation (short) skewness (short)	83.33 (10/12)	47.06 (8/17)
Nighttime, Sand, Long & Short	standard deviation (long) mean (short)	75.00 (9/12)	23.08 (3/13)

Table. 4.1. Feature selection and image and detection results using only the long and short wavelength images

Several conclusions can be drawn based on the results of this work. First, in general it is easier to detect "mines" buried in clay than it is to detect "mines" buried in sand. Second, there does not appear to be any clear advantage to using nighttime images over daytime images when detecting mines. Lastly, there seems to be little difference between results acquired using amplitude features of the long and short wavelength images and results acquired using amplitude features of the temperature and emissivity images.

Time of Day, Terrain, Image Pair	"Optimal" Feature Subset	Probability of Detection	Probability of False Alarm
Daytime, Clay, Temp & Emis	mean (temp) kurtosis (temp) skewness (emis) entropy (emis)	83.33 (15/18)	0.00 (0/15)
Nighttime, Clay, Temp & Emis	mean (temp) standard deviation (temp)	94.44 (17/18)	0.00 (0/17)
Daytime, Sand, Temp & Emis	mean (temp) standard deviation (temp) mean (emis) standard deviation (emis)	66.67 (8/12)	42.86 (6/14)
Nighttime, Sand, Temp & Emis	mean (temp) standard deviation (temp) skewness (temp) skewness (emis) kurtosis (emis)	91.67 (11/12)	7.64 (1/12)

Table. 4.2. Feature selection and image and detection results using only the temperatures and emissivity images

5. Supervised Learning Results in a Cluttered Environment

The data for this experiment was acquired from the NTS site. This site, which was located in the Nevada desert, was uncleared and therefore contained a large amount of clutter. Additionally, the mines were buried in a manner such that they were not visually apparent.

The clutter in these images, was due mainly to the presence of bushes. These areas had pixel intensity values which were much greater than the background and the mines areas. By using an automatic thresholding technique with the long wavelength image, the pixel locations of the clutter could be determined, and eliminated from further consideration as a mine. Since the ground truth of the fiducials was known, these areas were also masked from further consideration. After these steps, only mines, surrogates, and background areas were left unmasked in the images. The background consisted of areas of relatively unvegetated ground similar to where mines were actually buried. The sub-images which were sampled from these images were of size 13. This size was selected since it was approximately the diameter of a mine. For each sub-image, we computed nine amplitude features, nine texture features for a distance of one, and nine texture features for a distance of two. The amplitude features consisted of the mean, standard deviation, skewness, kurtosis, energy, entropy, local minimum, local maximum, and median. The texture features consisted of mean, standard deviation, contrast, angular second moment, correlation, entropy, local homogeneity, cluster shade for sum, and cluster prominence for sum. Each physical location consisted of 108 features, 27 features times 4 images (long, short, temperature, and emissivity). It is important to note that these results were acquired by using the long, short, temperature, and emissivity images together. This differs from the experiment described in section four, which analyzed the long and short wavelength images separately from the temperature and emissivity images.

Because this research is currently in progress, only nighttime images which contained metal mines and surrogates have been processed. These results are presented below.

5.1. Feature Selection

The feature selection was accomplished using 1480 background samples, 60 mine samples, and 45 surrogate samples. Since preliminary results indicated that the mine and surrogates had similar characteristics, they were combined into one class. The sequential forward selection algorithm was used to determine the "optimal" subset of features. Preliminary results indicated that by using the best three features as determined from the sequential forward selection algorithm, adequate separation between the two classes could be achieved. These three features in order of importance are the contrast of the emissivity for a distance of two, the local minimum value of emissivity, and the standard deviation of the emissivity for a distance one. The physical significance of these particular features seems to be related to the relative homogeneity of the soil surface where the holes were dug. Specifically, the soil surface is more homogeneous where the holes were dug. The differences, however, were subtle, as indicated by the probability density functions of the three features shown in Fig. 5.1. The image labeling using these features is described in the next section.

Figure 5.1. Probability density functions of the three "best" features. The smoothest curves represent the background statistics. The others represent mine and mine surrogate statistics.

5.2. Image Labeling

Feature vectors were generated for every unmasked pixel in the testing images. These feature vectors were then classified by a probabilistic neural network. The output was an image where each pixel represents the probability of a mine existing at that location. Next, this probability image could be thresholded to determine the most likely mine and surrogate pixels. The results of these operations are shown in Fig's 5.2 to 5.4. Fig. 5.2 displays the long wavelength, short wavelength, temperature, and emissivity images for a plot. Fig. 5.3 displays the probability image which results after the classification. The brighter areas of this image indicate locations for which a high probability of a mine exists. Lastly, Fig. 5.4 displays the labeled image which results from thresholding the probability image. This figure indicates the 500, 1000, 1500, and 2000 most likely mine pixels. In all these images, the squares indicate actual mine/mine surrogate locations. Because this research is currently in progress, post-processing results will not be discussed.

Figure 5.2. Long wavelength (a), short wavelength (b), temperature(c), and emissivity (d) images. Locations of mines and mine surrogate are indicated by squares.

Figure 5.3. Image representing the probability of a mine or surrogate pixel. The black area were masked out as they were identified as vegetation.

Figure 5.4. Most likely mine pixels : most likely 500 pixels (a), 1000 (b), 1500 (c), and 2000 (d) pixels are shown in white. Squares indicate mine positions.

6. Discussion

Using supervised learning pattern recognition techniques we found that we could detect land mines buried in an uncluttered clay soil from fused dual-band IR images with the use of simple amplitude features. While this exercise was instructive, the data did not require such sophisticated methods to detect mines. In fact, it was not clear that a two sensor IR system was necessary or even beneficial for the purpose. We did determine that probabilistic neural nets provided performance superior to nearest neighbor methods and even backpropagation neural nets. Using PNN we could distinguish between holes with and without mines.

However, further work with cluttered fields provided considerable instruction in to the utility of data fusion techniques. Current results indicate that the best performance was achieved using texture features extracted from the pseudo-image "emissivity". We currently believe this to be physically related to the surface disturbance resulting from the mine emplacement. We further expect that large scale shape features calculated from the original images will also add to the detection capability.

It is clear that even with these results, additional sensors must be included for reasonable detection reliability to be achieved. This data is currently being collected and will be included in the processing during the coming year.

7. Acknowledgment

"Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under contract number W-7405-ENG-48."

9. References

1. Ground penetrating radar data was provided in part by SRI, Inc.. Contact Roger Vickers (1-415-859-4422)
2. UV and hyperspectral visible/near IR imagery provided by SAIC, Inc.. Contact Michael Griggs (1-619-458-5200)
3. G. Wolberg, Digital Image Warping, IEEE Computer Society Press, 1990.
4. W.K. Pratt, Digital Image Processing, 2nd Edition, Wiley, pp.559-561.
5. R.M. Welch, K. Kuo, S.K. Sengupta, "Cloud and Surface Textural Features in Polar Regions", IEEE Trans. Geoscience and Remote Sensing, vol. 28, No. 4, pp. 520-528, July 1990.
6. T.Y. Young and K.S. Fu, Handbook of Pattern Recognition and Image Processing, Academic Press, 1986.
7. R.P. Lippmann, "An Introduction to Computing with Neural Nets," IEEE ASSP Magazine, April, 1987, pp. 4-22.
8. D.E. Specht, "Probabilistic Neural Networks," Neural Networks, Vol. 3, pp.~109-118, 1990.
9. N.K. Del Grande, G.A. Clark, P.F. Durbin, D.J. Fields, J.E. Hernandez, and R.J. Sherwood, "Buried Object Remote Detection Technology for Law Enforcement," SPIE Orlando '91 Symposium, Orlando, Florida, April 1-5, 1991.
10. G.A. Clark, J.E. Hernandez, N.K. Del Grande, R.J. Sherwood, S-Y Lu, and P.F. Durbin, "Computer Vision for Locating Buried Objects," Twenty-Fifth Annual Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, California, November 4-6, 1991.
11. "Airborne Detection of Buried Minefields," Energy and Technology Review, Lawrence Livermore National Laboratory, December, 1991.
12. J.E. Hernandez, Shin-Yee Lu, R.J. Sherwood, G.A. Clark, and B.S. Lawver, "A Signal and Image Processing Object-Based System Using CLOS," Lawrence Livermore National Laboratory Report. UCRL-JC-108409, October 28-November 1, 1991.
13. Hu, M.K., 1962. "Visual Pattern Recognition by Moment Invariants", IRE Trans. Info. Theory, vol. IT-8, pp. 179-187.
14. Devijver, P. A. and J. Kittler, 1982. Pattern Recognition: A Statistical Approach. Prentice Hall, Englewood Cliffs, 448 pp.
15. Narendra, P. and K. Fukunaga, 1977. "A Branch and Bound Algorithm for Feature Subset Selection", IEEE Trans. Comp., vol. C-26, no. 9, Sept. 1977, pp 917-922.
16. Specht, D., 1990b. "PNN and Polynomial Adaline as Complementary Techniques for Classification", IEEE Trans. Neural Networks, vol. 1, pp 111-121.
17. Cacoullos, T., 1966. "Estimation of a Multivariate Density", Annals of the Institute of Statistical Mathematics (Tokyo), vol.18(2), pp. 179-198.

18. Murthy, V.K. 1965. "Estimation of Probability Density", *Annals of Mathematical Statistics*, vol. 36, pp. 1027-1031.
19. Murthy, V. K., 1966. Nonparametric Estimation of Multivariate Densities with Applications, in P.R. Krishnaiah (Ed.), *Multivariate Analysis* (pp. 43-58), New York: Academic Press.
20. Parzen, E., 1962. "On Estimation of a Probability Density Function and Mode", *Annals of Mathematical Statistics*, vol. 33, pp. 1065-1076.
21. Rumelhart, D.E., J.L. McClelland and the PDP Research Group, 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, I & II*, Cambridge, MA: MIT Pr. 611 pp.
22. Hand, D. J., 1981. *Discrimination and Classification*, New York, L. Wiley and Sons, pp. 218.
23. A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice Hall, 1989.
24. Gonzales, Rafael C. and Wintz, Paul., *Digital Image Processing*, second edition. Addison-Wesley Publishing Company, 1987.
25. Hall, Ernest L., *Computer Image Processing and Recognition*. Academic Press, 1979.
26. Rosenfeld, Azriel and Kak, Avinash C, *Digital Picture Processing.*, Academic Press, 1976.
27. Wolberg, George, *Digital Image Warping*, IEEE Computer Society Press, 1990.
28. J. E. Hernandez, "Detecting and Locating Buried Mines from Dual-Band IR Data: A Statistical Pattern Recognition Approach", Lawrence Livermore National Laboratory Report (in preparation).
29. Table 2

END

DATE

FILMED

3/17/94