Conf-930445-22

UCRL-JC-113727
PREPRINT

RECEIVED
AUG 09 1993
OSTI

Sensor Feature Fusion for Detecting Buried Objects

G.A. Clark, S.K. Sengupta, R.J. Sherwood, J.E. Hernandez,
M.R. Buhl, P.C. Schaich, R.J. Kane, M.J. Barth, and N.K. DelGrande

This paper was prepared for submittal to the
SPIE International Society for Optical Engineering 1993 Symposium
April 12-16, 1993
Orlando, Florida

April 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.

MASTER

ds
DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

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 product, 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.

Sensor Feature Fusion for Detecting Buried Objects

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

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

Given multiple registered images of the earth's surface from dual-band infrared sensors, our system fuses information from the sensors to reduce the effects of clutter and improve the ability to detect buried or surface target sites. The sensor suite currently includes two infrared sensors (5 micron and 10 micron wavelengths) and one ground penetrating radar (GPR) of the wide-band pulsed synthetic aperture type.

We use a supervised learning pattern recognition approach to detect 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 subimage. 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, 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 extract features from the images, and use feature selection algorithms to select only the most important features according to their contribution to correct detections. This allows us to save computational complexity and determine which of the sensors add value to the detection system. The most important features from the various sensors are fused using supervised learning pattern classifiers (including neural networks). We present results of experiments to detect buried land mines from real data, and evaluate the usefulness of fusing feature information from multiple sensor types, including dual-band infrared and ground penetrating radar. The novelty of the work lies mostly in the combination of the algorithms and their application to the very important and currently unsolved operational problem of detecting buried land mines from an airborne standoff platform.

1. INTRODUCTION

The goal of this work is to detect and locate buried and surface objects, given multiple registered images of regions of the earth, obtained from a suite of various remote sensors. Past research has shown that it is extremely difficult to distinguish objects of interest from background clutter in images obtained from a single sensor. It is hypothesized, however, that information fused from a suite of various sensors is likely to provide better detection reliability, because the suite of sensors measures a variety of physical properties that are more separable in feature space. Materials surrounding the objects of interest can include natural materials (soil, rocks, foliage, water, holes made by animals and natural processes, etc.) and artifacts (objects made of metal, plastic and other materials). The sensor suite currently includes two infrared sensors (5 micron and 10 micron wavelengths) and one ground penetrating radar (GPR) of the short pulse, wideband synthetic aperture type. 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 uses 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. At the current time, our data set is limited, in that we have a small sample size. Our experience with GPR is preliminary at this time, so this paper focuses mainly on the fusion of images from dual-band infrared sensors.

2. EXPERIMENTS AND MEASUREMENTS

Our data were measured at the Lawrence Livermore National Laboratory. The soil is California adobe clay, densely populated with assorted sizes of gravel. In April, 1992, eighteen roughly identical, 1-foot diameter holes were dug in three 16 ft. by 16 ft. clay pits. One of two kinds of objects (plastic-cased mine surrogates and metal-cased mine surrogates) was placed at the same depth in each of nine holes. One additional pit contained surface mines. All holes were refilled. Tailings or spoils consisting of small clay clods, and 1/2 to 1-1/2 inch rocks surround some of the holes. The test pits are viewed by infrared sensors mounted on a 40-foot tower adjacent to the pits. The eighteen holes are visible on site, but the IR image resolution prevents human visual identification of the precise hole locations. The IR images of the mines are often much larger than the actual holes, because they indicate the disturbed soil from the tailings or adjacent ground.

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 is depicted in Fig. 1 and consists of four main parts: preprocessing, feature extraction, feature selection, and classification. These parts are used in a two step process to classify a subimage. 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, 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) [10].

3.1. Image Preprocessing

3.1.1 Acquisition

The dual-band infrared (DBIR) data were acquired using an Agema dual-band infrared scanner. The image data were stored real-time in the field on a 1.2 gigabyte disk which is part of the Agema system. After the experiment was completed, typical frames at appropriate times were selected, and written to 3.5 inch diskettes. The file format of these diskettes is compatible with the IBM PS/2, and the data format within the file is in an Agema-specific format. These diskettes were read on a Sun SPARCstation 2.

The images are stored by the Agema as four interlaced fields. We developed software on the Sun to reassemble the four fields into a single frame. Each image file from the Agema was passed through the software to generate a binary image file on the Sun in a format compatible with VISION, the image processing and machine vision package developed at LLNL. Images from the Agema are 140 pixels wide by 280 pixels high. We interpolated those images horizontally to 280 pixels wide by 280 pixels high.

3.1.2 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 [1]. 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 the same points 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 [1,22-24]. 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.

Figure 1. Fusion/Automatic Target Recognition Depend Heavily Upon Proper Image Representation.

3.1.3 Normalization

The images are normalized with respect to the background by subtracting the mean of the background from the images and dividing this result by the standard deviation of the background. This normalization aids in computing some of the image features and it makes the classifier less sensitive to absolute units, which can vary greatly with physical properties of the site from image to image. The estimation methods are detailed in [25].

3.1.4 Tile Cutting

We use ground truth information about the test site to manually cut out tiles (sub-images) containing data from mines, filled holes, clutter, background, fiducial markers, and surface mines. These sub-images then become the samples used for pattern classification.

3.2. Feature Extraction

Each image is divided into $N \times N$ pixel subregions centered at target pixels and background pixels. The value of the integer N for our problem varied from 20 to 40 pixels, but the most effective results were obtained for $N = 20$. We chose the tile (subregion) size so as to be contained entirely inside the boundary of the mine image (approximately), and not contain a significant amount of background information. For the IR images, approximately 20 pixels correspond to a foot in space.

Given preprocessed sub-images, we compute a vector of statistical features from the pixel values in the sub-images. Typical features include amplitude histogram features and texture features [2, 3]. Currently, we use only the amplitude histogram features (mean, standard deviation, skewness, kurtosis, energy and entropy). Future work will include use of spatially-dependent features, including texture features.

3.3. Feature Selection

Feature reduction 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-dimension 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. We gain two main benefits from this approach. 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 [3,4]. Commonly used algorithms include Branch and Bound, Sequential Forward Selection, and Sequential Backward Selection [3,4].

For our initial problem, we chose to use probability of detection as the feature selection criterion. Our feature vectors are of length twelve, because we use six histogram features per sensor and we use two IR sensors. This means there are only $4095 = 2^{12} - 1$ possible feature subsets, so we used a "brute force" search of all combinations to see which subset provided the maximum probability of detection (see section 4).

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 [11-20]. "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 [7-9, 25]. 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 a preprocessing 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 [4], the back propagation neural network [5], and the probabilistic neural network [6, 14]. Because we have a small sample size, we use the "hold one out" method for training and testing [4, 12].

4. SUPERVISED LEARNING RESULTS - MINES VS. BACKGROUND

In this experiment, we define a two-class problem. The *first class*, (called "Mines") corresponds to buried mines and filled holes. Both have similar IR characteristics, because IR detects surface effects, and both have disturbed soil. The *second*

class (called "No Mines") corresponds to background, clutter, fiducial markers, and surface mines (only buried mines were of interest in these tests). Image tiles were hand selected as described in section 3.0. Additionally, daytime and nighttime images were processed in an attempt to determine whether the time of day affected the probability of detecting mines.

For the daytime images, 24 "mine" samples and 178 "no mine" samples were used. For the nighttime images, 21 "mine" samples and 172 "no mine" samples were used. The samples are of size 21 pixels by 21 pixels. These samples were then run through the feature selection and classification process using the "hold one out" technique, and the probability of detection was computed for each of the 4095 possible feature sets. Table 1 summarizes the results. The probability of detection is high and the probability of false alarm is low, indicating very successful classification. The night-time measurements provided better performance, as expected. Previous research [7-9] shows that night-time IR measurements generally have higher signal-to-noise ratio than daytime measurements.

Table 1: Feature selection results.

Time of Day	"Optical" feature set	Probability of detection	Probability of false alarm
Day	stand. dev. (long) skewness (long) kurtosis (long) mean (short)	91.67 (22/24)	1.12 (2/178)
Night	mean (long) stand. dev. (long)	100.0 (21.21)	0.00 (0.172)

4.1 Image Labeling Results

The next step is to perform the classification using the features specified in the previous section over the entire image. For this process, a feature vector was generated for subimages centered at every pixel in the image and classified. The result is an image with "mine" and "no mine" pixels marked. The images produced by this process are shown in Figs. 2 and 3 for daytime and nighttime clay pit images, respectively. These figures show the results of processing only one image, but it should be noted that all the images produced similar results. The noticeable circular areas of the image which appear at regular intervals are fiducial marks which were used in the image pre-processing.

These results appear very favorable, especially for the nighttime images. The classifier detected all the mines accurately. The results for the daytime image are also very good, however slight inaccuracies do exist. One reason for these may be the fact that there tends to be more thermal clutter during the daytime than at night.

4.2. Image Labeling Results – Postprocessing

We see from Figures 3 and 4 that there exist some misclassified regions in the labeled image. This issue is resolved by applying a region growing algorithm [2] and size constraints to the images. The region growing algorithm connects similar pixels in neighborhoods and creates regions for analysis. Once we have identified regions, we apply a constraint (threshold) on the number of pixels in a region that can correspond reasonably to the number of pixels in a mine image. Figures 4 and 5 show the results of this post-processing step for the images corresponding to those in Figures 2 and 3. The bottom images show the results of applying region growing and size constraints to the images in Figures 2 and 3. In the top images in Figures 4 and 5, the shaded circles represent the ground truth for the mine locations. The light colored "dots" represent the results of our processing to estimate the centroid of the mine location. Clearly, the postprocessing reduces the number of false alarms and improves the reliability of the detections. Other postprocessing methods, including morphological operations, are available for dealing with connected regions, if necessary [2,4,21].

Fig. 2 Daytime IR images of the clay pit used in Section 4. The bottom image is from the long wavelength (10 micron) sensor. The center image is from the short wavelength (5 micron) sensor. The squares indicate ground truth for the locations of buried metal mines and filled holes. The noticeable circular areas which appear at regular intervals are fiducial markers used for image registration. The uppermost image is the labeled image, showing regions classified as a mine. The vertical and horizontal lines are artifacts from the scanning process used to publish the images in this paper.

Fig. 3 Nighttime IR images of the clay pit used in Section 4. The bottom image is from the long wavelength (10 micron) sensor. The center image is from the short wavelength (5 micron) sensor. The squares indicate ground truth for the location of buried metal mines and filled holes. The noticeable circular areas which appear at regular intervals are fiducial markers used for image registration. The uppermost image is the labeled image, showing regions classified as a mine. The vertical and horizontal lines are artifacts from the scanning process used to publish the images in this paper.

Fig. 4. Postprocessed labeled images for the Daytime IR measurements of the clay pit used in Section 4. The bottom image shows the results of applying region growing and size constraints to the image in Figure 2. In the top image, the shaded circles represent the ground truth for the mine locations. The light colored "dots" represent the results of processing to estimate the centroids of the mine locations.

Fig. 5. Postprocessed labeled images for the Nighttime IR measurements of the clay pit used in Section 4. The bottom image shows the results of applying region growing and size constraints to the image in Figure 3. In the top image, the shaded circles represent the ground truth for the mine locations. The light colored "dots" represent the results of processing to estimate the centroids of the mine locations.

5. SUPERVISED LEARNING RESULTS - MULTIPLE CLASS PROBLEMS

In this experiment, we investigated multiple-class problems, in which the classes are defined as follows: *Class BG*: Background, *Class O*: Filled hole, *Class 1*: Plastic Mine, *Class 2*: Metal Mine, and *Class M*: Mixed, plastic or metal mine (union of the two sets). One concern might be that the classifier might tend to see all objects as either a hole or not a hole. If that were true, then its discrimination capability would be limited. To test our concerns, we designed computer experiments in which we ask the classifier to detect the differences among the classes listed above. Figure 6 shows the IR images used in the experiments.

Fig. 6. IR images of three clay pits used in three- and four-class PNN and kNN classification. The dark squares mark the locations of the nine buried mines. The nine other light, bright, circular spots are filled holes. The dark spots in the corners of images are the fiducials used to register images. The other smaller dark marks arranged in geometrical patterns are markers. The images are day, long wavelength images. Short wavelength images look similar, as do night images.

Our studies indicate that given images from a single IR sensor, mines and filled holes cannot be separated by human eye or by a neural network. The goal of this test is to see if the multiple classes can be separated using two IR bands and two different classifiers; the probabilistic neural network (PNN) and the k-nearest neighbor (kNN) classifier (where $k = 1$). The PNN converges to the Bayes optimal classifier as the sample size becomes large. The kNN classifier is simple, accurate and commonly used, so it provides us with a benchmark for comparison. We also used a back propagation neural network (BPNN), but we did not show its results, because in many cases, it did not converge well.

We used the processing method described in section 3.0, except that after doing feature selection, we decided to use all of the histogram features for both IR bands. The results are summarized in Table 2. The PNN performs very well, without the false

minima problem of the BPNN. PNN also outperforms kNN, in general, for this problem. Both PNN and kNN outperform human visual inspection.

	PNN						NN						
Tile kind	BG	0	1	2	M	Pd	BG	0	1	2	M	Pd	Tile size
BG	47	1	0	0	—	.98	47	1	0	0	—	.98	40 x
0	0	16	1	1	—	.89	1	12	3	2	—	.67	
1	0	0	6	0	—	1.0	0	4	2	0	—	.33	
2	0	0	0	12	—	1.0	0	2	0	10	—	.83	
BG	47	1	—	—	0	.98	47	1	—	—	0	.98	40
0	0	16	—	—	2	.89	1	13	—	—	4	.72	
M	0	2	—	—	16	.89	0	4	—	—	14	.78	
BG	48	0	0	0	—	1.0	48	0	0	0	—	1.0	20 x
0	1	11	3	3	—	.61	1	11	3	3	—	.67	
1	1	4	1	1	—	.67	0	3	2	1	—	.33	
2	0	1	0	11	—	.92	0	2	0	10	—	.83	
BG	48	0	—	—	0	1.0	48	0	—	—	0	1.0	20
0	1	12	—	—	5	.67	0	13	—	—	4	.72	
M	0	2	—	—	16	.89	0	6	—	—	12	.67	

5.1. Observations

(1) The four class problem shown in Table 2 indicates that given dual-band IR measurements, the classifiers can detect the differences among plastic mines, metal mines and holes. (2) Comparing the four class problem with the three class problem, we see that, as expected, when the plastic and metal mines are mixed (three class problem) the classifiers have more difficulty separating the classes than when the classes contain clearly different objects (four class problem).

6. GROUND PENETRATING RADAR

We have preliminary results with ground penetrating radar (GPR) of the short-pulse, wideband, synthetic aperture radar (SAR) type, operating in the 500 MHz to 1 GHz range with two dipole antennas. It viewed our test pits from about 10 ft. elevation and -45 degrees look angle from horizontal. SRI International (SRI) provided the required GPR system and operated it for the tests performed at LLNL's test pits.

The SRI radar system uses a high voltage pulse generator for the transmitter excitation and a high speed data acquisition system for the receiver. Additional circuit elements are used for bandlimiting the received signal, and dynamically altering the receiver gain. Two separate antennas are used for the transmit and receive functions and consist of resistively-loaded dipole radiators with corner reflectors to improve directivity. The imaging system uses a software-based quadrature demodulation system and conventional SAR processing.

During the tests at LLNL the system was configured to generate a pulse every 0.25 and 0.5 in of linear travel. An 80 kW pulse was generated of which 10-20 kW was in the 500-1000 MHz frequency band. The antenna configuration was positioned for vertically polarized transmit and receive electric field vectors.

The vehicle-mounted radar system was moved along a straight line parallel to one boundary of the test pit area. The antenna system was mounted at a height of approximately 10 ft off the ground and a view angle of -45 degrees from the horizontal was obtained for the test pits.

SAR images of the buried land mine pits were produced from the data collected at LLNL. Initial results did not provide adequate resolution of the targets due to a lack of proper 'focusing' of the SAR image. The image quality was improved through additional processing and the final result was images with 6-12 in resolution. Figure 7 shows the final result of the image processing using this hardware.

A number of artifacts exist within the final images: a 'ghost' image is present in front of, and behind, the strong targets, and some displacement of the targets is also present. The ghost images are most likely due to the broad pulse shape generated by the transmitter and antenna. Displacement of the targets is due to the change in rf propagation velocity for electric fields in the ground. Some additional position errors may have been introduced through tilt of the vehicle and antennas as it moved along the path adjacent to the test area. Tilt of the vehicle would result in a range variation that would not be accounted for in the images. Future plans include developing methods to correct for these artifacts.

Fig. 7. SAR Image Data of Mine Pit. The soil is clay, and the pit contains three metal mines and three filled holes. The squares on the image depict the ground truth positions of the mines and holes. The crosses depict the locations of fiducial markers. The metal mines are located at the corners of an imaginary triangle with its peak at the top of the image. The holes are located in the remaining positions (see Figures 2 and 3).

As expected, the GPR images show strong signatures of metal mines, but very weak or nonexistent signatures of plastic mines. The sample size is extremely small, so we choose to not report the specifics of the preliminary detection results at this time. However, we are encouraged by the fact that for metal mines in some scenarios (e.g. metal mines in sand), while the IR images are weak and indistinguishable by eye from background and clutter, the GPR images are very strong. We are therefore hopeful that the ATR system using fused IR and GPR images will outperform the system using IR images only.

7. DISCUSSION

Supervised learning pattern recognition techniques perform well in detecting land mines buried in clay soil from fused dual-band IR images. Future work includes the acquisition of a much larger data set, including more IR and GPR measurements. We also plan to further incorporate image segmentation algorithms and rule-based processing into the preprocessing step, and use spatial features, such as those that characterize texture. Real-time implementations and system integration for airborne platforms will follow.

Two other areas for future investigation are important. Correction of the GPR image by accounting for the change in propagation medium parameters should improve position estimates and potentially the object detection probability. Use of the received signal phase history may also contribute to the ability to detect the presence of a buried object provided the amplitude response indicates something of interest in the immediate area.

A second area that requires investigation is an analysis of the hardware and its ability to resolve differences in return signal amplitudes. The digitizer systems employed have been 8-bit converters which limit the received signal dynamic range to approximately 48 dB. The return signal from a plastic mine may or may not be detectable with this dynamic range limit. The goal of this investigation is to determine how much difference is required in material properties for the detection of a plastic mine element.

8. ACKNOWLEDGMENT

We gratefully acknowledge the support of David Fields, Michael Carter, Mark Eckart and Arthur Toor of LLNL, and the Defense Advanced Research Projects Agency.

"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. G. Wolberg, *Digital Image Warping*, IEEE Computer Society Press, 1990.
2. W.K. Pratt, *Digital Image Processing*, 2nd Edition, Wiley, pp.559-561.
3. 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.
4. T.Y. Young and K.S. Fu, *Handbook of Pattern Recognition and Image Processing*, Academic Press, 1986.
5. R.P. Lippmann, "An Introduction to Computing with Neural Nets," *IEEE ASSP Magazine*, April, 1987, pp. 4-22.
6. D.E. Specht, "Probabilistic Neural Networks," *Neural Networks*, Vol. 3, pp. 109-118, 1990.
7. 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.
8. 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.
9. "Airborne Detection of Buried Minefields," *Energy and Technology Review*, Lawrence Livermore National Laboratory, December, 1991.
10. 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 rpt. UCRL-JC-108409, October 28-November 1, 1991. `end{thebibliography}`
11. Hu, M.K., 1962. Visual Pattern Recognition by Moment Invariants, *IRE Trans. Info. Theory*, vol. IT-8, pp. 179-187.
12. Devijver, P. A. and J. Kittler, 1982. *Pattern Recognition: A Statistical Approach*. Prentice Hall, Englewood Cliffs.
13. 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.
14. Specht, D., 1990b. PNN and Polynomial Adaline as Complementary Techniques for Classification, *IEEE Trans. Neural Networks*, vol. 1, pp 111-121.
15. Cacoullos, T., 1966. Estimation of a Multivariate Density, *Annals of the Institute of Statistical Mathematics (Tokyo)*, vol.18(2), pp. 179-198.
16. Murthy, V.K. 1965. Estimation of Probability Density, *Annals of Mathematical Statistics*, vol. 36, pp. 1027-1031.
17. 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.
18. Parzen, E., 1962. On Estimation of a Probability Density Function and Mode, *Annals of Mathematical Statistics*, vol. 33, pp. 1065-1076.
19. 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.
20. Hand, D. J., 1981. *Discrimination and Classification*, New York, L. Wiley and Sons, pp. 218.
21. A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice Hall, 1989.
22. Gonzales, Rafael C. and Wintz, Paul., *Digital Image Processing*, second edition. Addison-Wesley Publishing Company, 1987.
23. Hall, Ernest L., *Computer Image Processing and Recognition*. Academic Press, 1979.
24. Rosenfeld, Azriel and Kak, Avinash C, *Digital Picture Processing*, Academic Press, 1976.
25. 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).

END

**DATE
FILMED**

10 / 28 / 93