Conf-9210231--4

UCRL-JC-112103
PREPRINT

Computer Vision and Sensor Fusion for Detecting Buried Objects

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

This paper was prepared for submittal to the
Asilomar '92 Systems, Signal, and Computers
Pacific Grove, California
October 26-28, 1992

October 1992

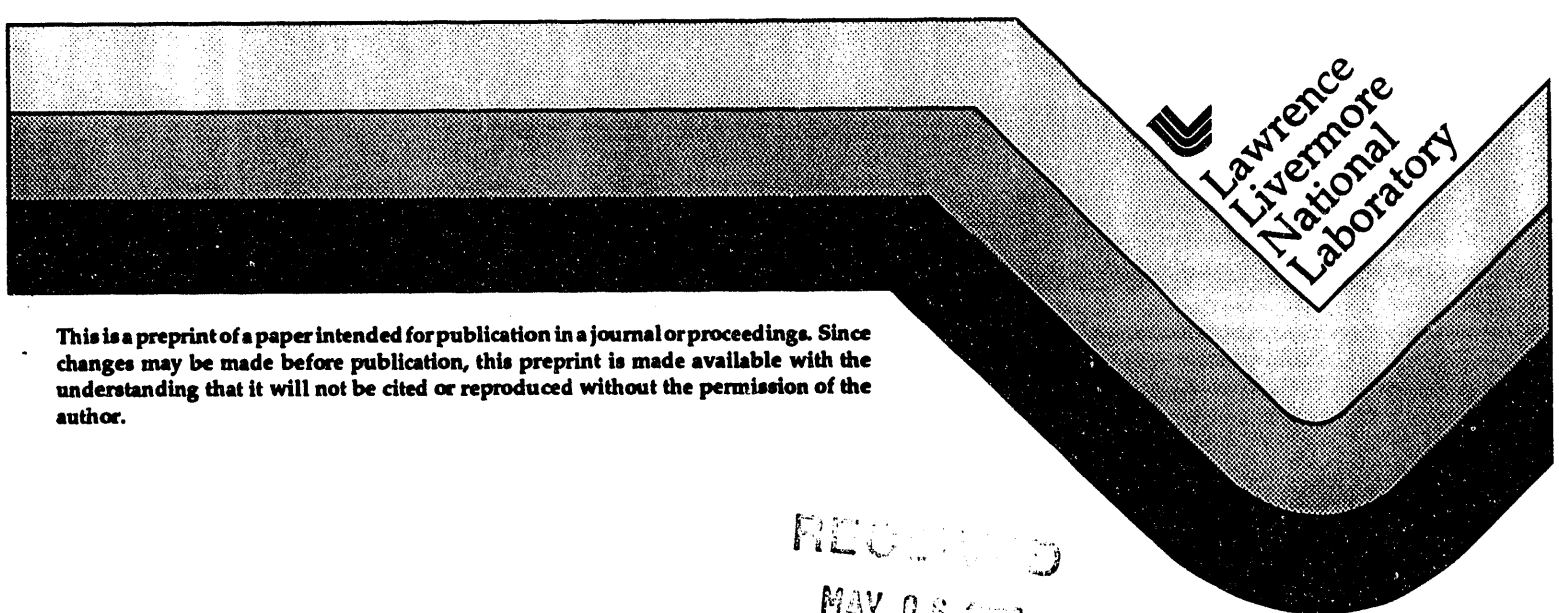

DISCLAIMER

Computer Vision and Sensor Fusion for Detecting Buried Objects

Gregory A. Clark, Jose E. Hernandez, Sailes K. Sengupta, Robert J. Sherwood,
Paul C. Schaich, Michael R. Buhl, Ronald J. Kane, and Nancy K. DelGrande
Lawrence Livermore National Laboratory
Electrical Engineering
Livermore, CA, 94550

Abstract

Given multiple images of the surface of the earth 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. Supervised learning pattern classifiers (including neural networks) are used. We present results of experiments to detect buried land mines from real data, and evaluate the usefulness of fusing information from multiple sensor types. The novelty of the work lies mostly in the combination of the algorithms and their application to the very important and currently unsolved 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 detects a variety of physical properties that are more separable in feature space. The 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 pulsed 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 mine detection. 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 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. Six months ago, 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 easily 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 between 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 Jose E. Hernandez). VISION is an object-oriented package, and it runs under Franz Allegro CL, which implements the Common Lisp Object System (CLOS).

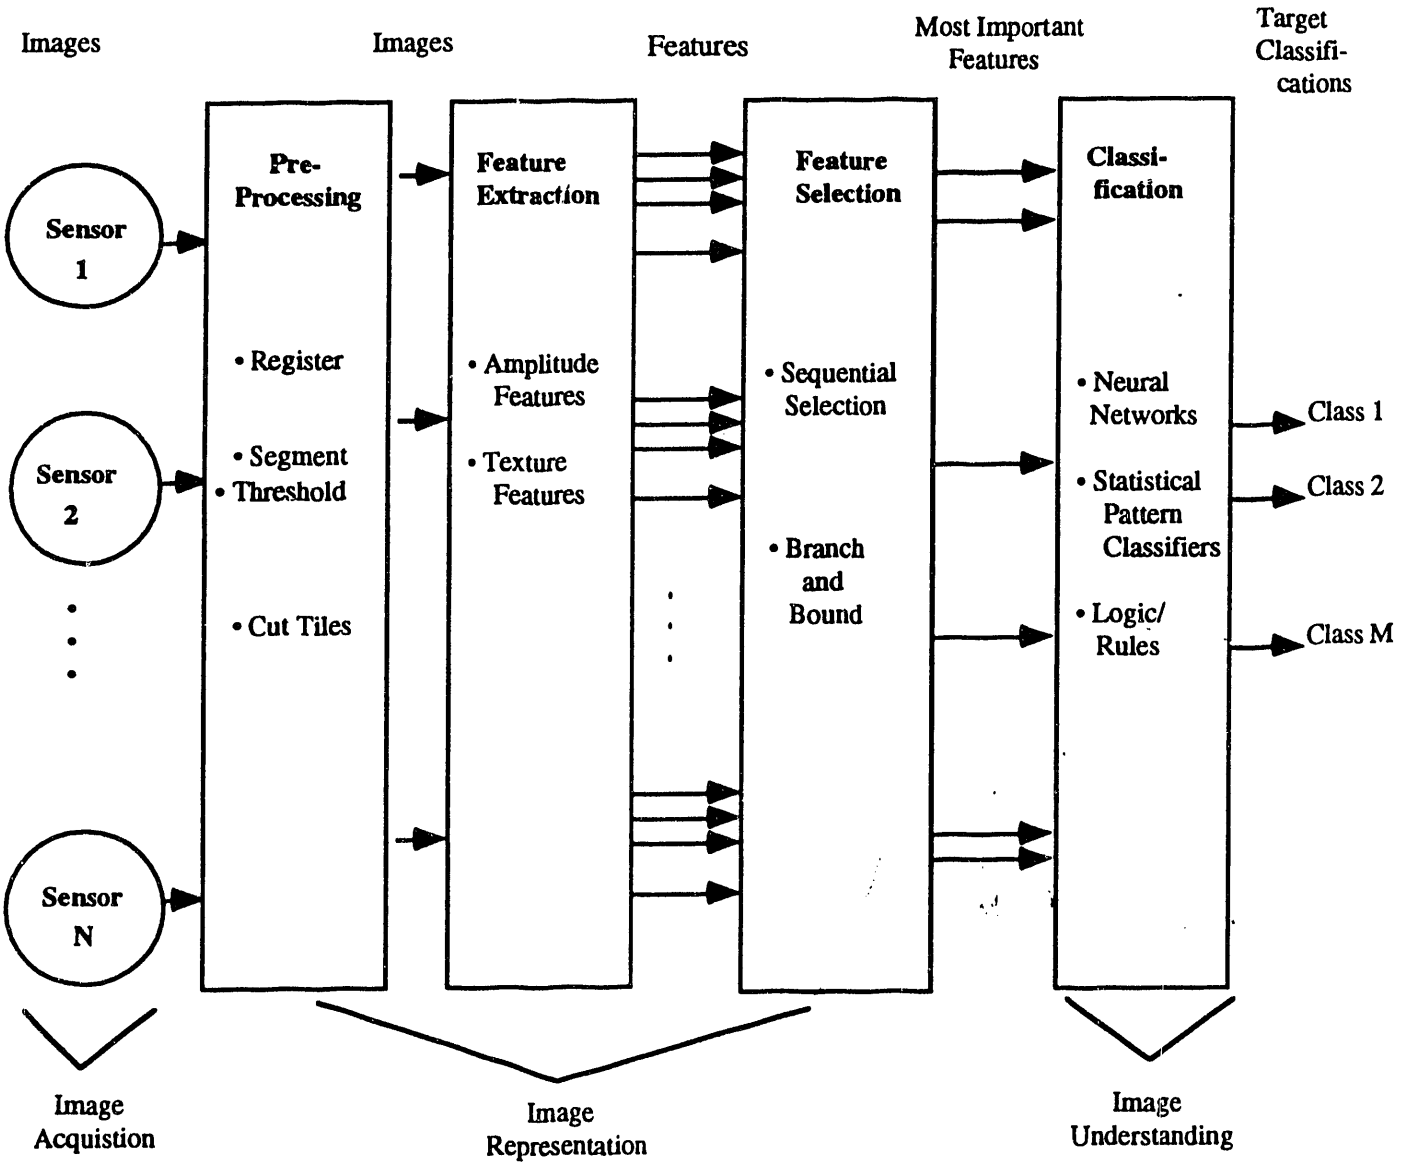

Figure 1: Fusion/Automatic Target Recognition Depend Heavily Upon Proper Image Representation.

3.1 Image Preprocessing

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 same point on the surface.

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.

Tile Cutting: We use ground truth information about the test site to manually cut out tiles containing sub-images of 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

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).

3.3 Feature Selection

An important goal is to select the subset of features that contribute most to correct classification. 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, 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, e 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 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.

3.4 Classification

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]. Because we have a small sample size, we use the "hold one out" method for training and testing.

Table 1: Features selection results.

Time of Day	"Optimal" 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 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 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.

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 images, respectively. These figures show the results of processing only one image, but it should be noted that all the images produced similar results. In these images, the square outline in the long and short wavelength images indicate where "mines" actually exist, and the white regions in the labeled image indicate were the classifier selected "mine" pixels. The other

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 good, especially for the night time 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. Fortunately, the mislabeled regions in the daytime images are small, and may be eliminated by a rule-based process (using shape and size constraints).

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 0:* 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 it's 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 4 shows the IR images used in the experiments.

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. 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 backpropagation 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.

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 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).

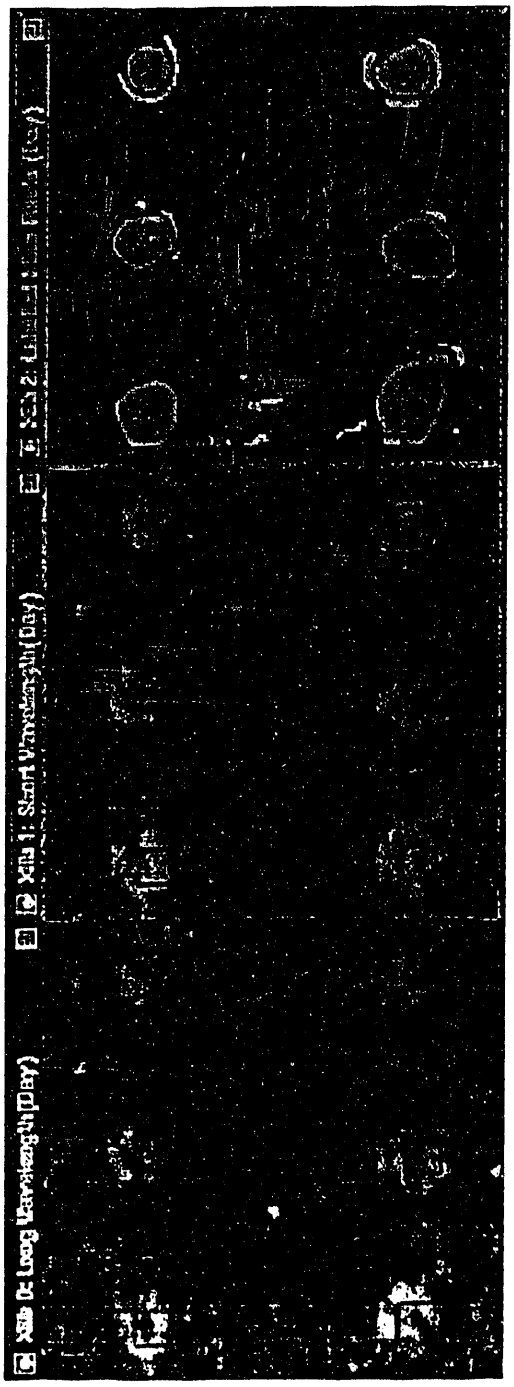

Figure 2: Daytime IR images of the clay pit used in Section 4. The bottom image is from the long wavelength (10 micron) sensor. The center image is from the short wavelength (5 micron) sensor. The squares indicate ground truth for the locations of buried metal mines and filled holes. The uppermost image 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.

Table 2: PNN and kNN performance on six features from dual-band IR imagery of Clay (see Fig. 4).

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

6 Ground Penetrating Radar and Other Sensors

We have preliminary results with ground penetrating radar (GPR) of the pulsed, synthetic aperture type, operating in the 200 MHz to 1 GHz range with two dipole antennas. It viewed our test pits from about 14 ft. elevation and 19 degrees look angle from horizontal. 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 results at this time. However, we are encouraged by the fact that for metal mines in some scenarios, 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. Future work will include investigating the use of surface measurements from visible, laser reflectance, near IR and UV sensors to help separate surface effects from the effects of buried objects.

7 Discussion

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

Acknowledgement

We gratefully acknowledge the support of David Fields, 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.

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, Pergamon Press, 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

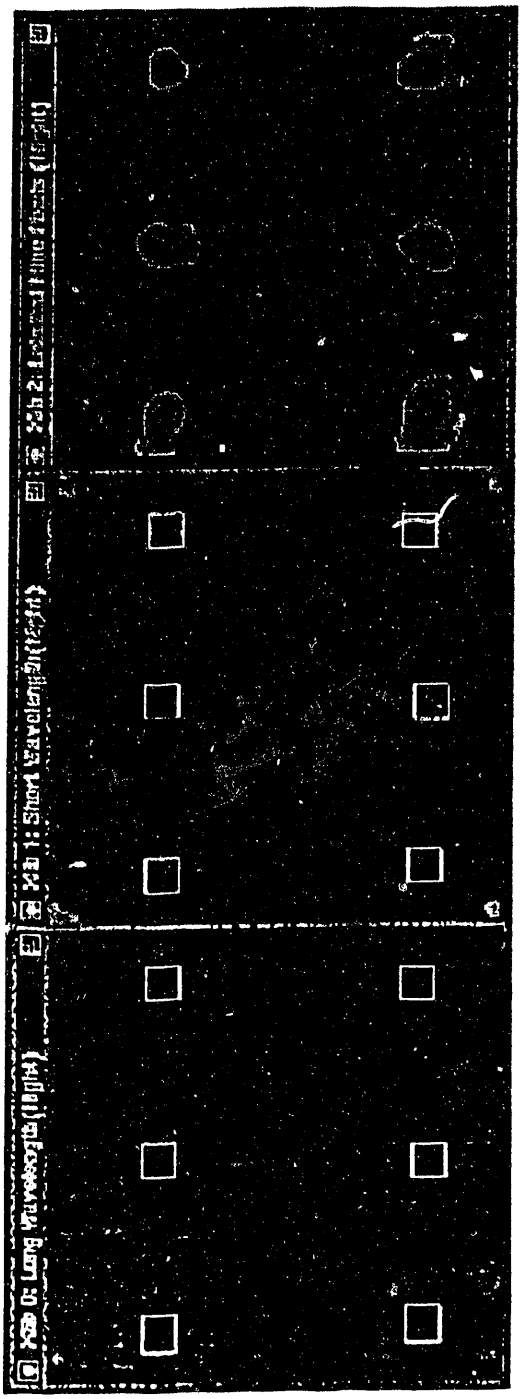

Figure 3: Nighttime IR images of the clay pit used in Section 4. The bottom image is from the long wavelength (10 micron) sensor. The center image is from the short wavelength (5 micron) sensor. The squares indicate ground truth for the location of buried metal mines and filled holes. The uppermost image is the labeled image, showing regions classified as a mire. The vertical and horizontal lines are artifacts from the scanning process used to publish the images in this paper.

Figure 4: 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.

CLOS," Lawrence Livermore National Laboratory report UCRL-JC-108409, presented at the First Annual LISP Users and Vendors Conference, Gaithersburg, Maryland, October 28-November 1, 1991.

END

DATE
FILMED
8/11/93