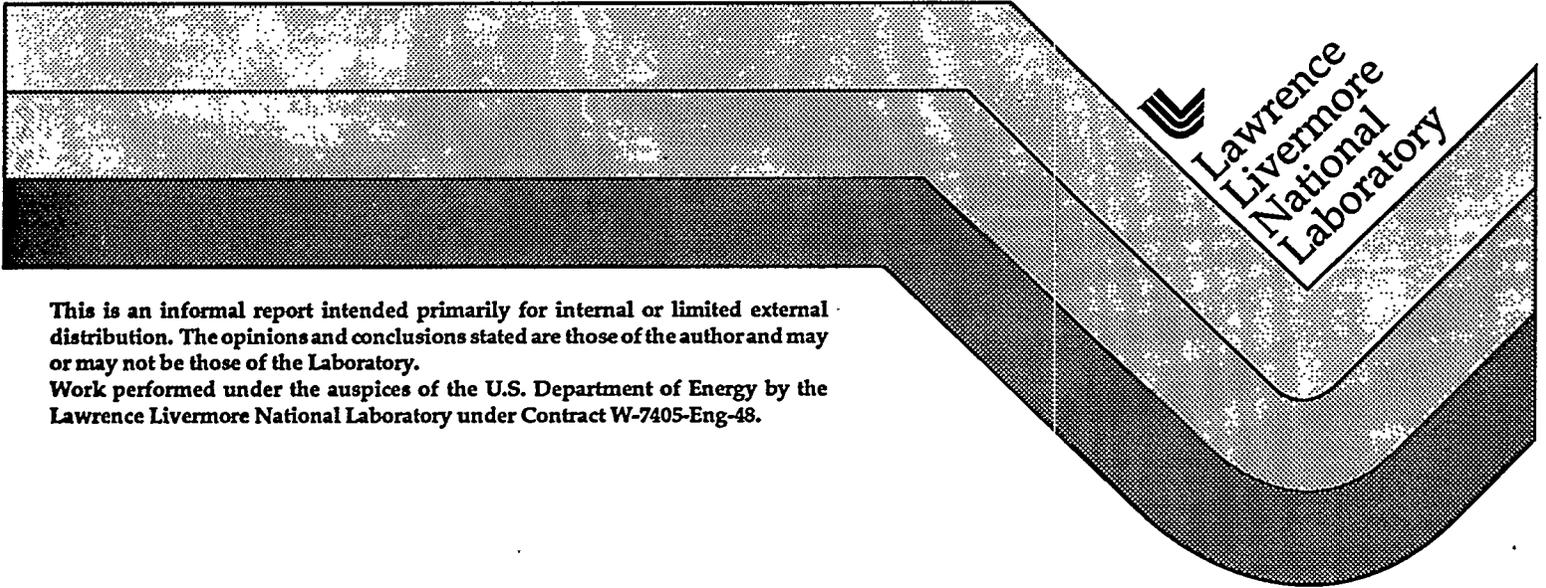


12
3-22-95 JS①

LLNL Electro-Optical Mine Detection Program

C. Anderson, W. Aimonetti, M. Barth, M. Buhl, N. Bull, M. Carter,
G. Clark, D. Fields, S. Fulkerson, R. Kane, C. Lee, F. Lee,
B. Mc Kinley, J. Page, F. Roeske, Jr., T. Rossow, P. Sargis,
J. Scarafiotti, P. Schaich, S. Sengupta, and R. Sherwood

September 30, 1994



This is an informal report intended primarily for internal or limited external distribution. The opinions and conclusions stated are those of the author and may or may not be those of the Laboratory.
Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make 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 any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

LLNL Electro-Optical Mine Detection Program

Charles A. Anderson, William D. Aimonetti, Marvin J. Barth, Michael R. Buhl, Nathan Bull, Michael R. Carter, Gregory A. Clark, David J. Fields, Stephen Fulker-son, Ronald J. Kane, Chuck K. Lee, F. Dean Lee, Billy J. Mc Kinley, James R. Page, Franklin Roeske Jr., Terry L. Rossow, Paul D. Sargis, John J. Scarafiotti, Paul C. Schaich, Sailes K. Sengupta, and Robert J. Sherwood

September 30, 1994

MASTER

Lawrence Livermore National Laboratory

AS
DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Table of Contents

List of Figures	iii
List of Tables	iv
Executive Summary	1
I. Introduction	3
History	3
Objectives	4
Program Activities	4
II. Test Site Development	5
III. Data Collections	13
Thermal Infrared Sensors	13
UV, Visible and Near-Infrared and Hyperspectral Sensors	16
Ground Penetrating Radar Sensors	17
Other Data Collections	19
IV. Data Processing	20
LLNL data acquisition and processing	20
Automatic Target Recognition and Data Fusion	24
Supervised Learning Results for a Clutter-Free Environment	26
Supervised Learning Results for a Cluttered Environment	30
Discussion	33
V. Summary of Results	34
Visual Inspection of data	34
Summary of ATR Results	39
VI. References	40
Appendix A- Ground-penetrating Radar for Buried Mine Detection	A1
Appendix B- Preliminary Results of LLNL Analysis of CSS COBRA Flight. .B1 Test Data for Mine Detection	
ENCLOSURE 1 - UV Mine Detection Field Test Data	
ENCLOSURE 2 - Hyperspectral Imaging of Emplaced Land Mines	
ENCLOSURE 3 - Ground-based Ultrawideband Synthetic Aperature Radar Measurements	

List of Figures

Figure 1. Test Facility at NTS	5
Figure 2. Aerial Photo of Controlled area.....	6
Figure 3. Map showing current location of mines and test plots	7
Figure 4. Motorhome and lift truck with sensors mounted on platform..	7
Figure 5a. M15, M19 mines and Czech surrogate before covering	9
Figure 5b Emplacement of Italian VS-1.6 mine	10
Figure 6. Composite visible image of J Southafter mine emplacement...	11
Figure 7. Diagram of mine and mine surrogate placement.....	12
Figure 8. Other test objects and targets	12
Figure 9. Infrared and Visible Cameras	14
Figure 10. Schematic Diagram of Dual Band Infrared System	15
Figure 11. Schematic Diagram of LLNL GPR System.....	18
Figure 12. Flow of DBIR data from capture to ATR processing.....	22
Figure 13. Flow of GPR data from acquisition to ATR processing.....	23
Figures14.1. Probability density functions of best features.....	31
Figure 14.2. SWIR, LWIR, temperature and emissivity images	32
Figure 14.3. Probability of a mine or surrogate pixel	32
Figure 14.3. Most likely mine pixels and mine locations.....	33
Figure 15. Visible image of sub-plot MS2	34
Figure 16. Infrared composite of plot J South.....	36
Figure 17. IR Images of sub-plots JS2 and JS3	37
Figure 18a. GPR Image of plot I South (Metal Mines)	38
Figures18b. GPR Image of plotK North (Plastic Mines)	38

List of Tables

Table 1. Mine Inventory.....	8
Table 2.1. Long and Short Wavelength Results.....	29
Table 2.2 .Temperature and Emissivity Results	29

LLNL Electro-Optical Mine Detection Program

Executive Summary

Under funding from the Advanced Research Projects Agency (ARPA) and the U.S. Marine Corps (U.S.M.C.), Lawrence Livermore National Laboratory (LLNL) has directed a program aimed at improving existing detection capabilities against buried mines and munitions. This program was begun in December of 1991 to serve as a center of effort for work in this area by ARPA and the services. Since that time the program has provided a national test facility for buried mines in arid environments, compiled and distributed an extensive data base of infrared (IR), ground penetrating radar (GPR), and other measurements made at that site, served as a host for other organizations wishing to make measurements, made considerable progress in the use of ground penetrating radar for mine detection, and worked on the difficult problem of sensor fusion as applied to buried mine detection. While the majority of our effort has been concentrated on the buried mine problem, LLNL has worked with the U.S.M.C. on surface mine problems as well, providing data and analysis to support the COBRA (Coastal Battlefield Reconnaissance and Analysis) program.

The original aim of the experimental aspect of the program was the utilization of multiband infrared approaches for the detection of buried mines. It was generally acknowledged that no single sensor would provide reliable capabilities over a broad range of conditions but it was considered possible to use multiple bandpasses to improve detection and reduce false alarms. Later the work was extended to a multisensor investigation, including sensors other than infrared imagers.

After an early series of measurements at a small site located at LLNL, it was determined that further progress in technology development and verification would require a much larger test facility which would provide investigators with a large sample of targets, preferably in a "natural" environment. The Buried Object Test Facility was constructed at the Nevada Test Site, north-west of Las Vegas, Nevada to serve this function. It is currently the largest stable test site for mine detection technology development, with 500 mine targets, including 300 live (unfused) mines of U.S. and overseas origin. In addition to being home to a large set of targets, the site is well documented, instrumented for both weather and ground temperature monitoring, and well maintained.

A number of measurements have been conducted at the NTS Buried Object Detection Facility, with sensors spanning the electromagnetic spectrum from the near ultraviolet (UV) to radio frequencies. The majority of the measurements were made with multispectral infrared, with separate images taken at 3-5 μ m and 8-12 μ m. This data was taken over the year from May 1993-June 1994. In addition color photographs and panchromatic images were taken of all sites. Hyperspectral visible data covering the near UV to the near infrared (NIR) was taken as well, but because of the enormous volumes of data produced by this method, it was only done for a subset of the field. Extensive GPR measurements were made of all targets in the field under a variety of conditions. The IR and GPR measurements have been carefully coregistered with ground truth and are available as well documented data sets. These were

made available to interested parties, and several organizations have used these data sets. In addition to measurements sponsored by this program, measurements have been made by other investigators. MIT/Lincoln Laboratory brought and made measurements with a stepped CW GPR rail-SAR system. ERIM made measurements of infrared spectra at the site. Neither of these data sets have been made available to us at the time of this writing.

While the instrument and measurement technologies for infrared is quite mature, the use of radar to detect or identify buried objects is not. Typically, radar systems are dragged slowly across the ground to make what are analagous to sounding measurements. Such a technique holds little promise for a practical wide area mine detection system, even for mine clearing and proofing. Recent work in impulse radar has suggested that airborne systems could be capable of detecting mines. Significant progress was made in demonstrating the imaging of mines using synthetic aperture radars (SAR) using impulse waveforms. While measurements were made from the ground, they can be extrapolated to airborne applications. In addition, ground based radars could be used for demining purposes.

Initial analysis of the infrared data indicated that single band IR data does not provide sufficient information to support a robust buried mine detection capability. In earlier investigations, the thermal signatures of mines in pristine, clutter free environments was sufficient for detection. While mine signatures were present in the single band data taken in a natural environment, the level of thermal clutter resulted in high rates of false alarms for any acceptable detection rate.

The use of two infrared bands (3-5 μ m and 8-12 μ m) improved the detectability of buried targets, verifying the initial premise that this two band data carried useful information on the mine signature. The results were, however, still not adequate for robust detection. The primary problem remains the inability to discriminate between mine signatures, when present, and the rather high level of natural clutter.

The ground penetrating radar system demonstrated a high probability of detection for buried metal mines even in this natural environment. It is clearly the most effective stand-off sensor for such targets. Imaging of plastic mines was also demonstrated, but, as in the infrared, the natural clutter resulted in an unacceptable false alarm rate.

At least five possible paths for improvement have been identified:

- 1) improved spatial resolution, which may provide better ground texture discrimination,
- 2) an analysis which involves more complicated spatial queueing and filtering,
- 3) additional IR bands, specifically using imaging spectroscopy,
- 4) the use of additional sensors other than IR imagers and the use of data fusion techniques with multi-sensor data, and
- 5) the use of time dependent observables such as temperature.

While these improvements present great challenges when applied to wartime scenarios, each of these paths are likely to be helpful in less constrained, post-hostility demining operations.

I. Introduction

This report summarizes Electro-Optical Mine Detection program activities conducted by Lawrence Livermore National Laboratory for FY92 through FY94. These activities were funded by the Defense Advanced Projects Research Agency and the U.S. Marine Corps as a part of a national effort to enhance the countermine capability of the armed services to detect mine fields at "standoff" distances. To this end, this effort focused on combining advances in multiple sensors, signal processing and sensor fusion toward establishing a technically feasible and tactically plausible approach to the countermine needs.

History

Since the 1970s, the mainstream countermine detection system developers and users conducting research on standoff mine field detection have been: the Defense Advanced Research Projects Agency (DARPA), the U.S. Army Belvoir Research, Development, and Engineering Center (BRDEC), the U.S. Army Engineering Waterways Experiment Station (WES), the U.S. Marine Corps Research, Development and Acquisition Command (MCRDAC), the U.S. Naval Coastal Systems Center (NCSC) presently called the Coastal Systems Station (CSS) and various elements of the user community. Their main efforts were focused on characterizing passive infrared (IR) and active laser (near-IR) signatures of surface mines.

In 1989 DARPA funded an LLNL program under the name of TEMPS (Temperature Evaluated Mine Position Survey). The TEMPS dual-band IR method used ratios of IR images (centered near 4.7 and 10.6 microns) to locate buried mine sites which typically had thermal and emissivity-related signatures unlike those associated with surface object clutter. In January 1991, TEMPS was successfully demonstrated using helicopter fly-overs of a Byron California grass-covered field with buried surrogate mines and clutter. This first successful proof-of-principle demonstration used separate temperature and emissivity-ratio maps and rule-based detection algorithms to verify all buried surrogate anti-tank mine site signatures and remove most of the clutter.

Additional LLNL funds were used to accelerate the DARPA-sponsored work, in response to the newly appreciated threat posed by mines during and after the Gulf War. In March 1991, helicopter fly-overs of the Yuma Arizona Army Proving Grounds successfully imaged the sites of all 36 live mines buried to a depth of 1 to 15 cm in dry sand. As part of an incipient coordinated countermine effort, LLNL received funding through DARPA for concept refinement of the dual-band infrared method for airborne standoff minefield detection. The plan was to evaluate advances in relevant technologies (such as multiple sensor techniques, image processing and sensor fusion) which would provide a more robust, all-weather system, with fewer false site detections, better able to distinguish buried mine sites from clutter.

During the same period, efforts by the U.S. Army Corps of Engineers Waterways

Experient Station (WES) in Vicksburg, Mississippi, the U.S. Navy NCSC in Panama City, Florida, and the U.S. Army Engineering BRDEC at Fort Belvoir Virginia continued. Advanced Technology Transition Demonstrations demonstrated the reliability of the military systems for detecting surface laid mine fields. Unfortunately, the effectiveness of these systems against the more difficult problem of buried mines is not yet considered satisfactory.

We also conducted a series of measurements with other investigators at a small site located at LLNL. Multi-sensor data were taken in April, 1992, with supplemental data taken at other times. The scenes imaged included a variety of mines at various depths of burial. In addition, images of some surface objects were taken in order to make initial evaluations of the potential for clutter rejection. The imagery was taken with ultraviolet (UV), visible, near-infrared (NIR)[1], and dual-band infrared (DBIR) passive cameras, an active IR camera, and a ground penetrating radar (GPR) system[2] processed with synthetic aperture radar (SAR) algorithms.

After completing these measurements and initial analysis it was determined that a larger test facility with actual mines was required to provide statistically significant evaluations of the multiple sensors and ATR algorithms in a realistic environment. This report summarizes program activities from October 1992 to October 1994. The analysis of the initial data taken at LLNL in 1992 is also discussed.

Objectives

The overall objective of the LLNL Electro-Optical Mine Detection program is to identify methods for standoff minefield detection, with specific focus on buried and surface scatterable mines.

The initial goals of the project were:

1. to build and instrument a mine test site at the NTS to gather UV, visible, near IR, dual-band thermal IR, and ultrawide band (200–2000 MHz) radar images;
2. to collect preliminary multisensor image data which allow direct comparisons and co-registration of UV, visible, near IR, dual-band thermal IR and radar images; and
3. to develop and apply the most appropriate automatic target recognition and clutter reduction algorithms to classify mine sites, increase detection probability and decrease false alarms for individual mines.

Program Activities

The principal activities of the LLNL Electro-Optical Mine Detection program have been: 1) the construction and maintenance of test facilities 2) the collection and coregistration of data taken with a set of sensors which represents a multi-spectral sensor suite; 3) the development and application of tools for data analysis to develop and evaluate new techniques for buried mine detection; and 4) preparations for follow on studies.

II. Test Site Development

A small clutter free site was established during 1992 at LLNL 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 surrogate. Care was taken to avoid leaving obvious surface signatures.

In 1993 LLNL established the test facility shown in figure 1 at the DOE Nevada Test Site (about 65 miles NW of Las Vegas, NV) to provide a permanent test facility

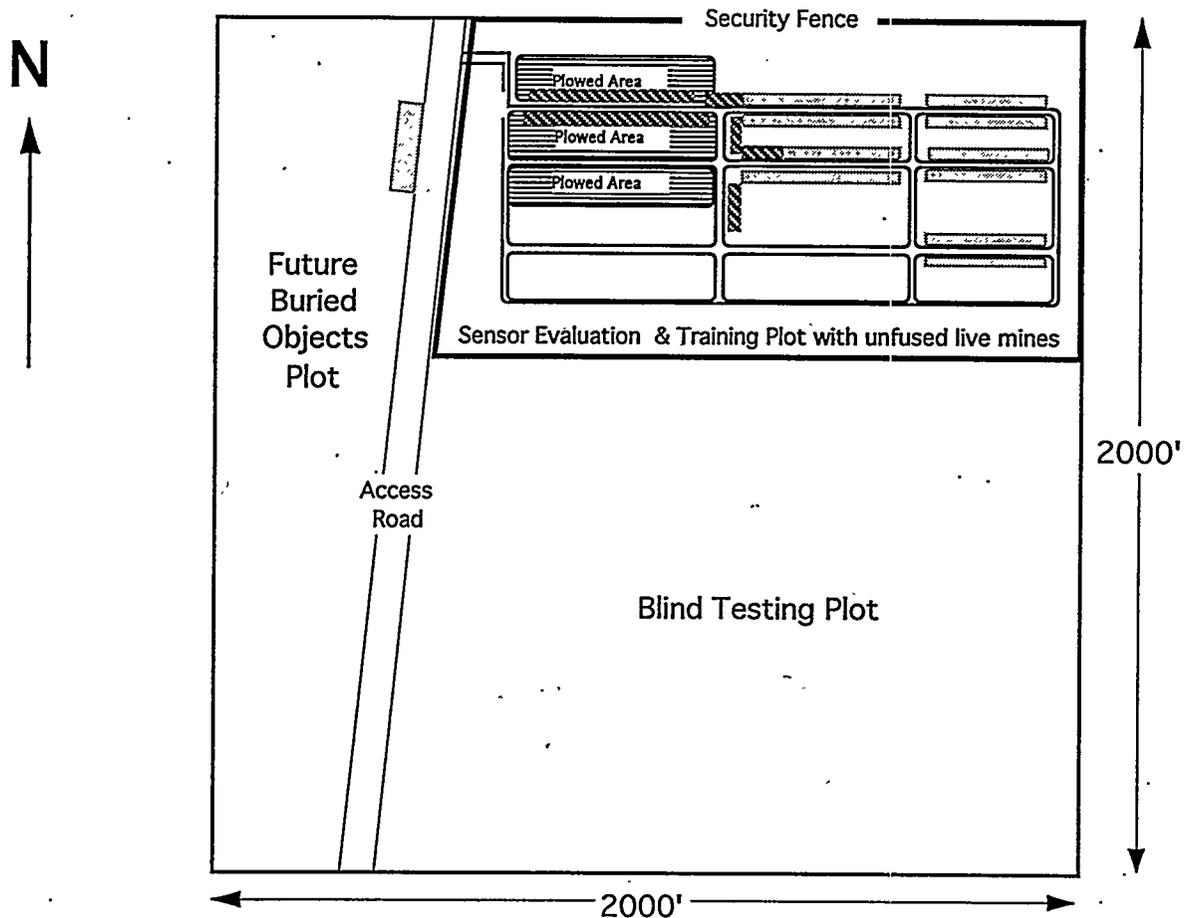


Figure 1. NTS Buried Objects Test Facility

with a large number of targets which would allow statistical evaluation of different technologies and ATR algorithms. The facility covers 100 acres, of which 30 acres are enclosed by a security fence with intrusion alarms to provide a controlled area (figure 2) for buried high explosives. The area is typical of the high desert basins of the western U.S. with modest to sparse natural vegetation and includes both cleared

and natural areas. The native soil is unconstituted alluvium; chiefly sand, silt, peb-

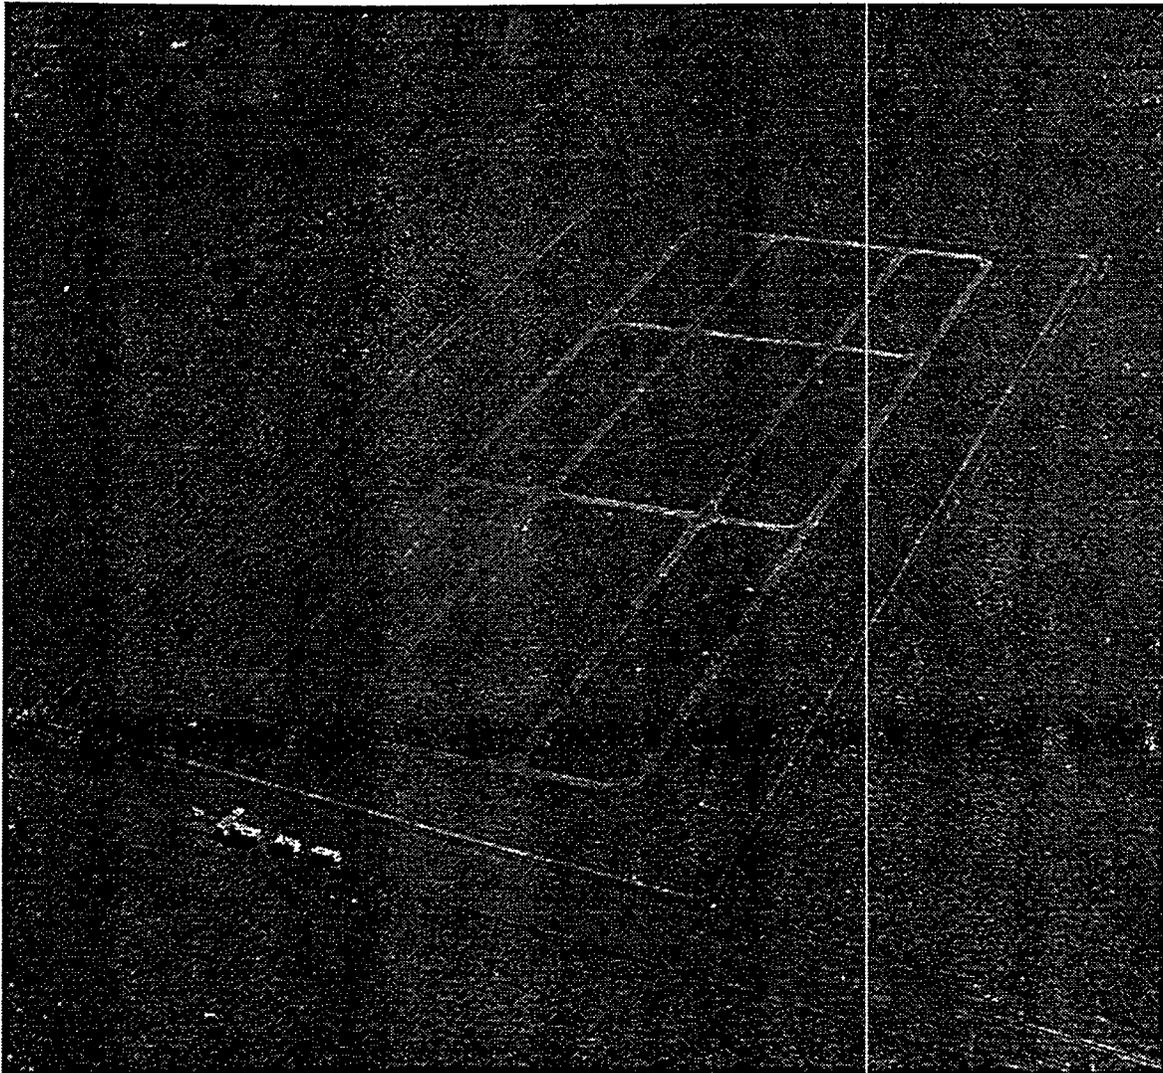


Figure 2. Aerial view of controlled area

bles and cobbles. The facility is in a low RFI area which provides a benign environment for radar and microwave sensor testing. The site was approved in February of 1993, and the mines were deployed in May of 1993. Different types of mines and other test objects were buried along the roads as shown in Figure 3 to allow either ground or airborne data collections. Storage trailers, electrical power, and telephones are located at the facility.

Data capture and data processing equipment were located in motorhomes and the sensors were mounted on a 60 foot platform lift with remote pan and tilt mount (Figure 4).

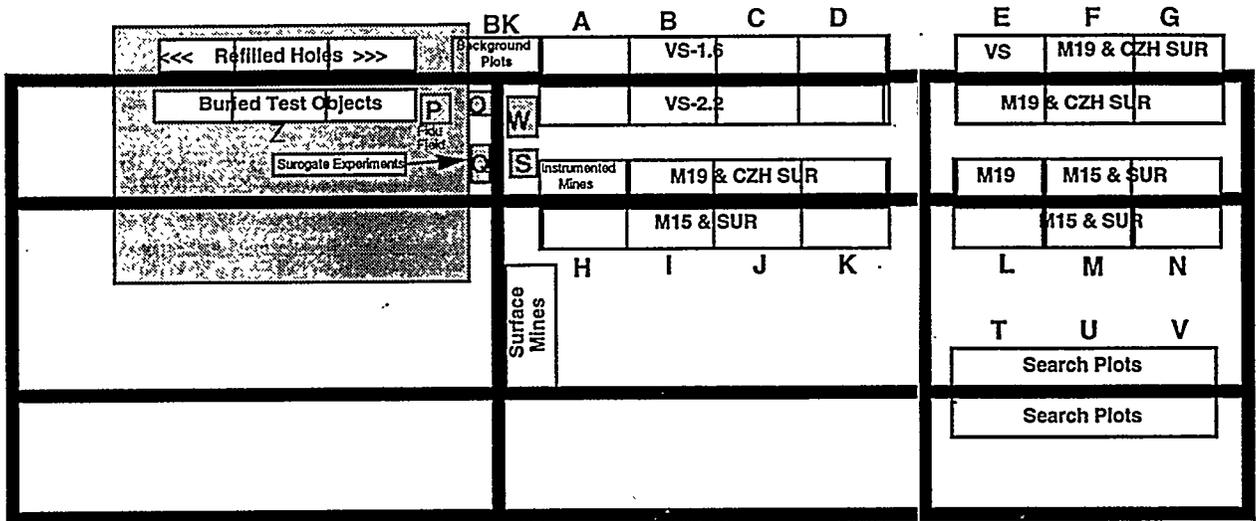


Figure 3. Location of Mines and Test Objects in controlled area

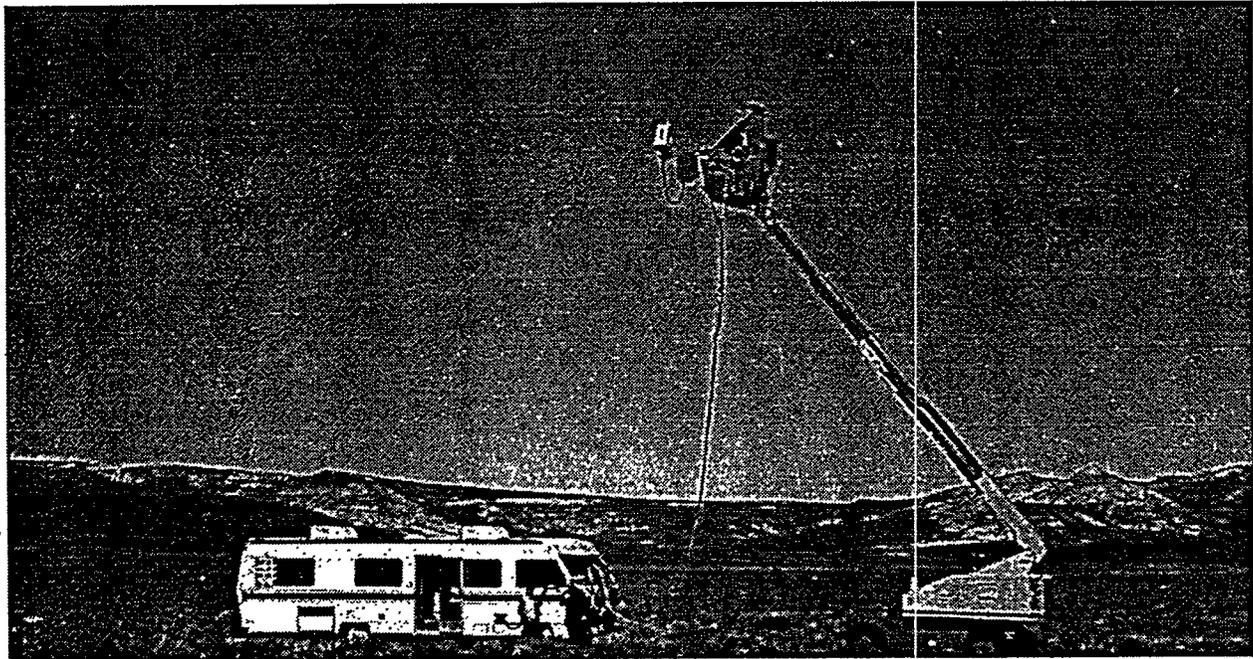


Figure 4. Instrumentation motorhome and lift truck in the controlled area

The current buried mine inventory includes the following:

Table 1:

Type	Description	Number
M-15	Metal - Live	100
M-15	Metal-Surrogate	100
M-19	Large Plastic - Live	100
M-19	Large Plastic - Surrogate	20
Czech	Large Plastic - Surrogate	100
VS-1.6	Medium Plastic - Live	50
VS-2.2	Medium Plastic - Live	50

Figure 5a shows the M-19 and M15 mines and the Czechoslovakian (PT-Mi-Ba-III) plastic surrogate before being covered. The surrogates for the plastic mines were filled with a grout mixture and provided a reasonable match to the thermal properties of actual mines but did not simulate the correct index of refraction for the plastic mines. Figure 5b shows the emplacement of the Italian VS-1.6 mines. The mines were buried carefully, between 0.5" and 4" deep, so that little or no surface signature was apparent to an observer (figure 5b and figure 6). Locations of the mines and fiducials for image registration have been accurately surveyed and are recorded in a ground truth database. The locations of the different mines in the controlled area are shown in figure 3. Figure 7 shows the planned locations of the mines in a typical training/test plot. The actual location of the mines in the plots were determined by the need to avoid bushes. A number of mines were randomly buried in the search plots, (T,U, V in figure 3) at locations unknown to experimenters to provide true evaluations of detection capabilities and false alarm rates. Background plots with no mines in them were also laid out with fiducials to allow registration of images of background plots. The target inventory also includes a few mines of different types placed on the surface. Other test objects (e.g. drums, pipes and plates) have also been buried as shown in figures 3 and 8.

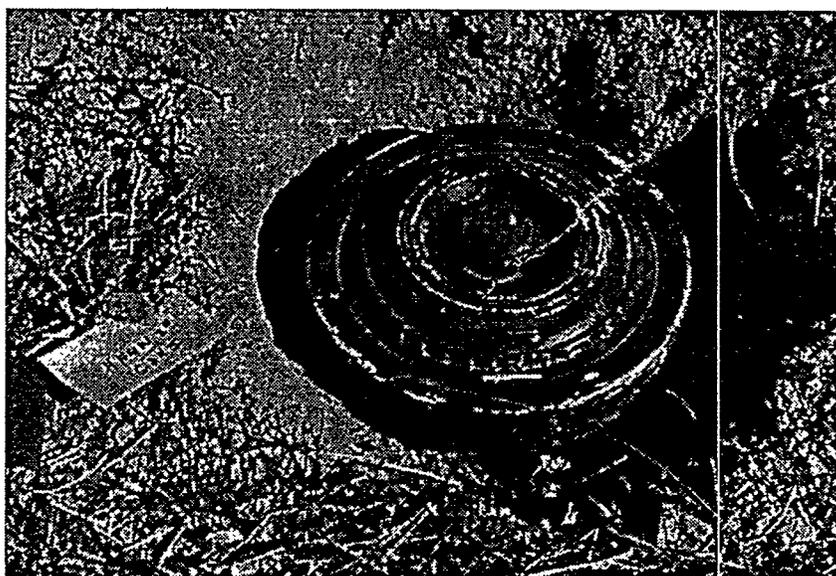
The test site was instrumented with a meteorological station which measured the air temperature, humidity, wind velocity, and net surface solar radiation flux. Temperature probes were also located at several distances above and below the ground surface. Several temperature sensors were also located on the surface above each type of mine and on its surface. Radiometer measurements at 4 microns and 10 microns of the ground surface temperature above instrumented mines and in background areas were made over a one month period with instruments provided by WES.



U.S. M15 Metal anti-tank mine



U.S. M19 Plastic anti-tank mine



Czech Plastic mine surrogate

Figure 5a. Emplaced mines before covering

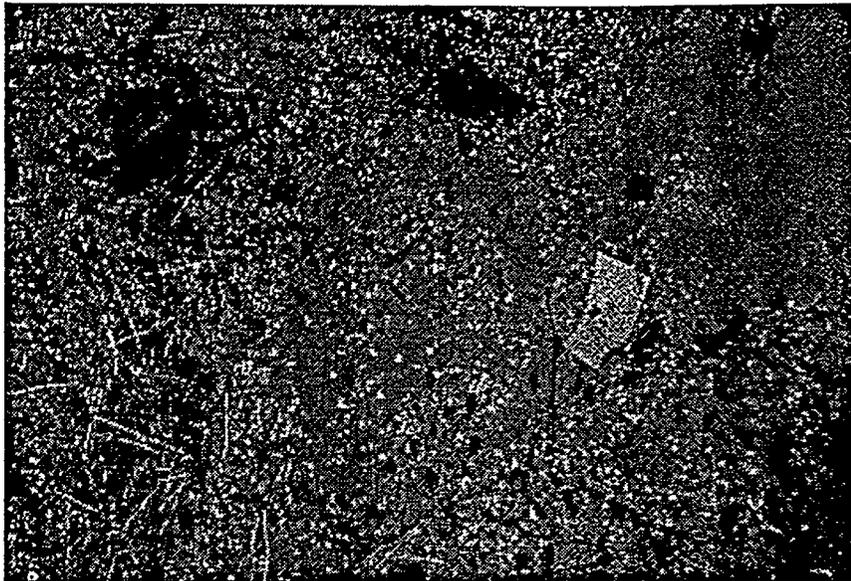
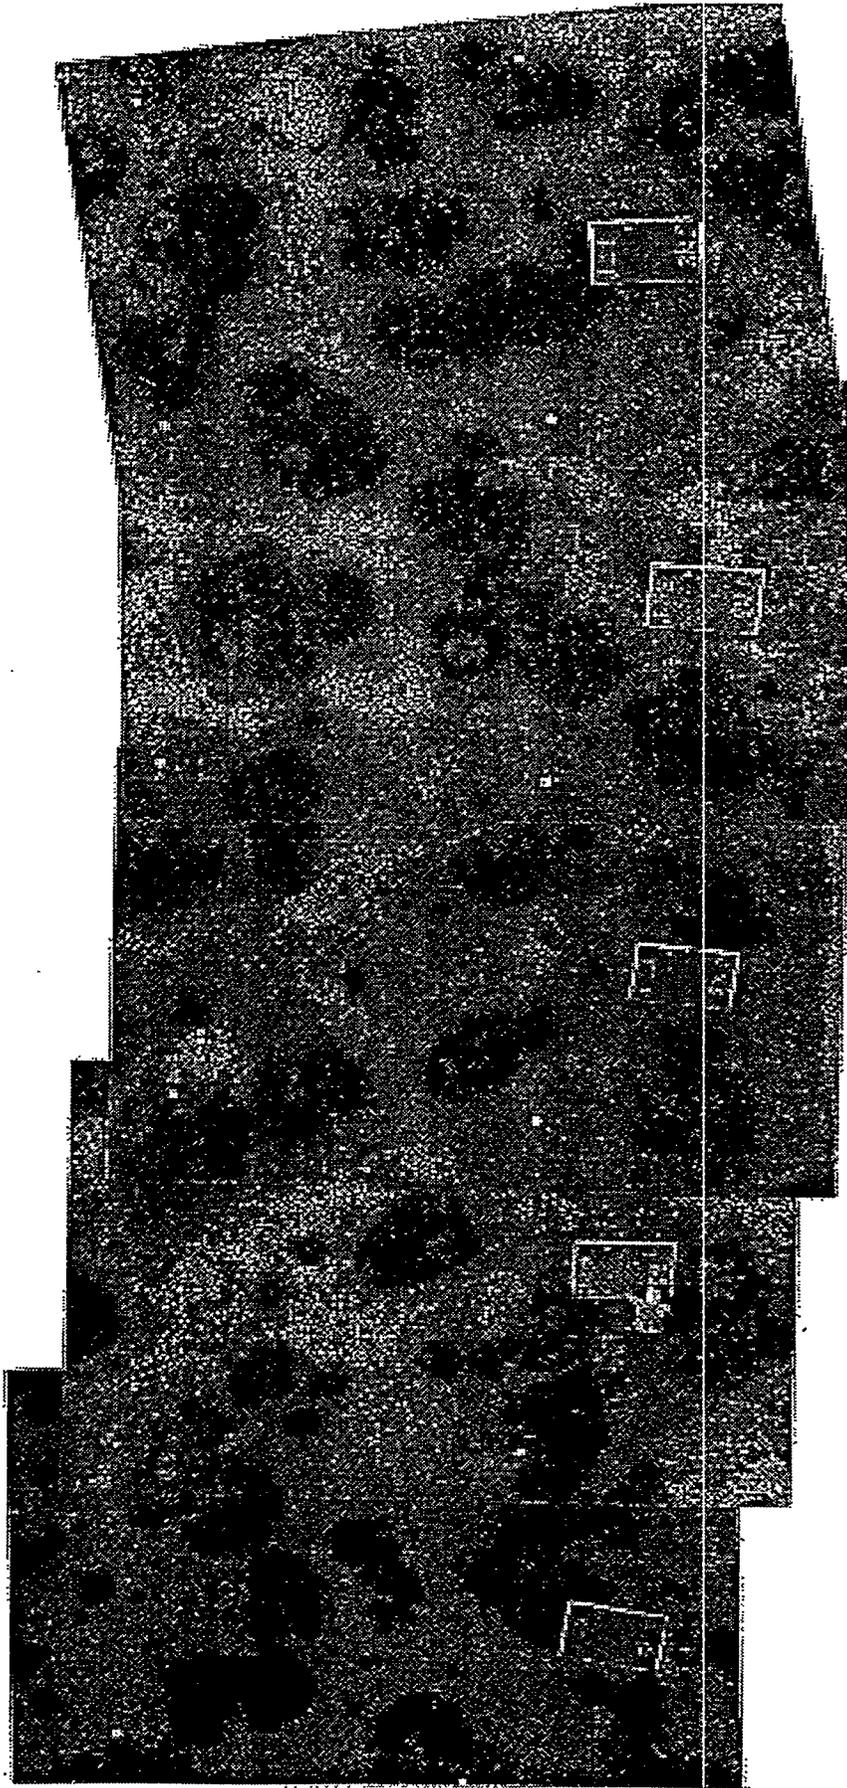


Figure 5b. Emplacement of Italian VS-1.6 plastic anti-tank mine



40120-688-01

Figure 6. Composite visible image of Plot J South after mine emplacement.
The white markers are 4" square metal fiducials for image registration.

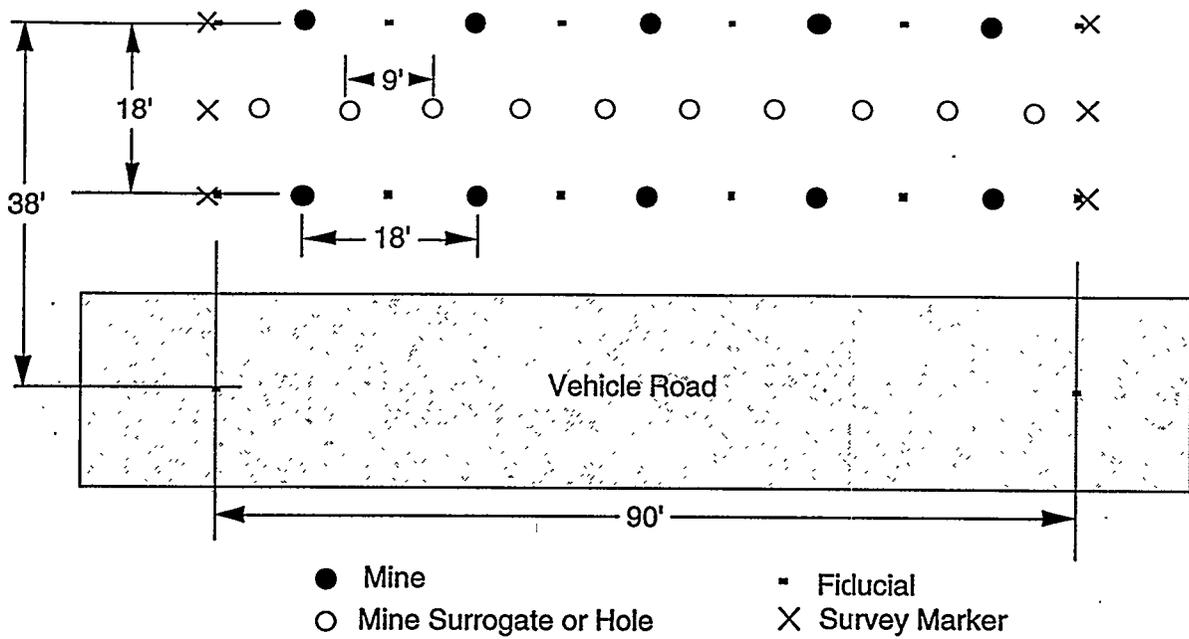


Figure 7. Mine and surrogate placement in typical plot

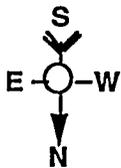
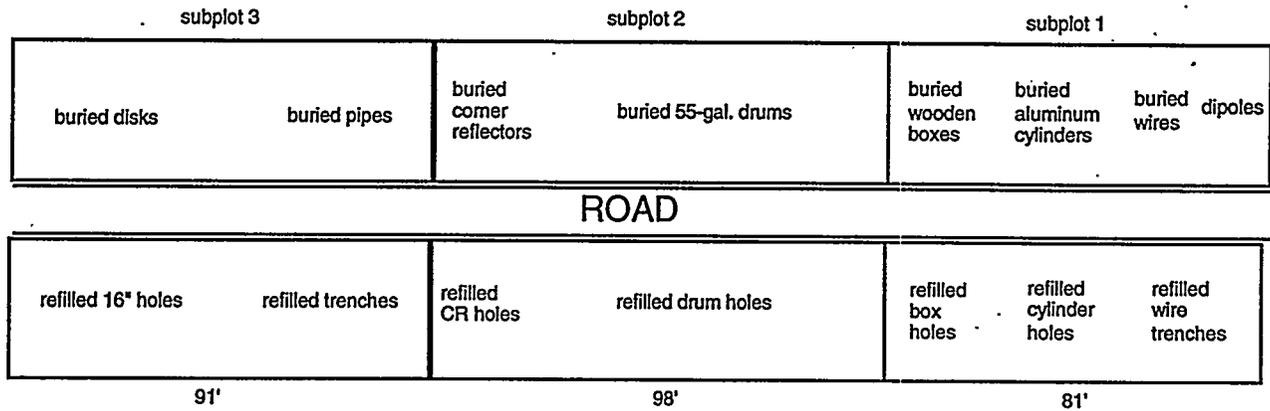


Figure 8. Other buried test objects

III. Data collections

We conducted a series of measurements with other investigators [1,2] at the small site located at LLNL. Multi-sensor data were taken in April, 1992, with supplemental data taken at other times.

A large number of data collections with a variety of sensors have been made in FY1994 at the NTS Facility. Dual band IR measurements were made of all mines at various times and times of year. LLNL and SRI[1, 31]) have made measurements with several configurations of a ground penetrating impulse radar system. Both color and monochrome daylight images were acquired. SAIC[2] acquired hyperspectral data sets and UV data sets on a sample of both buried and surface mines. These measurements are available to interested investigators. Two other contractors have visited the site to make measurement of their own. MIT/Lincoln Laboratories brought a stepped frequency-CW radar rail-SAR to the site and made measurements on a few of the mines. ERIM used a infra-red spectrometer to make measurements of reflected spectra from the ground surface from the mines and background regions. We do not have reports or data from either MIT or ERIM.

Thermal Infrared Sensors

Physical principles

Thermal IR images provide a measure of the apparent temperatures of the surfaces they image. These images are partly emitted surface radiation, which relates to the surface temperature, and partly reflected sky radiation which relates to the cooler sky temperature. Corrections must also be made for absorption and emissions of the air between the ground surface and the sensor. The apparent surface temperature equals the true surface temperature when the surface emissivity is equal to one (i.e., there is no reflected IR component) and the surface is referred to as an ideal blackbody surface. Most surfaces differ from a blackbody both by having non-zero reflectivity, and in having a spectral dependence of this quantity.

A single broad-band spectral measurement is insufficient to separate image data related to the surface material temperatures from image data related to the surface emissivity or reflectance properties. Because of the differences in material properties, the radiant emission measured in two or more wavelength bands provides information that can help separate these two quantities. Use of these techniques may allow one to distinguish between the native surface and superficial objects.

If the emissivity of the surfaces in a scene are uniform, the differences are due to differences in surface temperature. The thermal transport properties of a buried mine, covered by disturbed surface materials, produce a surface temperature above the mine which differs from the ambient value. Buried mine sites are typically warmer by day and cooler by night than undisturbed soil sites. While buried mine sites may have different surface reflectance properties (shortly after burial) than their sur-

roundings, some or all of this difference usually disappears after weathering. Thus the normal signature of buried mines is a result of a combination of thermal effects and altered surface properties.

Instrumentation

The infrared (IR) thermal images were recorded using an Agema 880 Dual-Band IR Camera System. This system had a short wavelength InSb detector with a filtered wavelength interval of 3.3 to 5.0 mm and a long wavelength Hg-Cd-Te detector with a filtered wavelength interval of 9.7 to 11.6 mm. 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 instruments are capable of resolving temperature differences equivalent to 0.1°C. The Agema cameras (figure 9) and computer were located on a platform lift truck

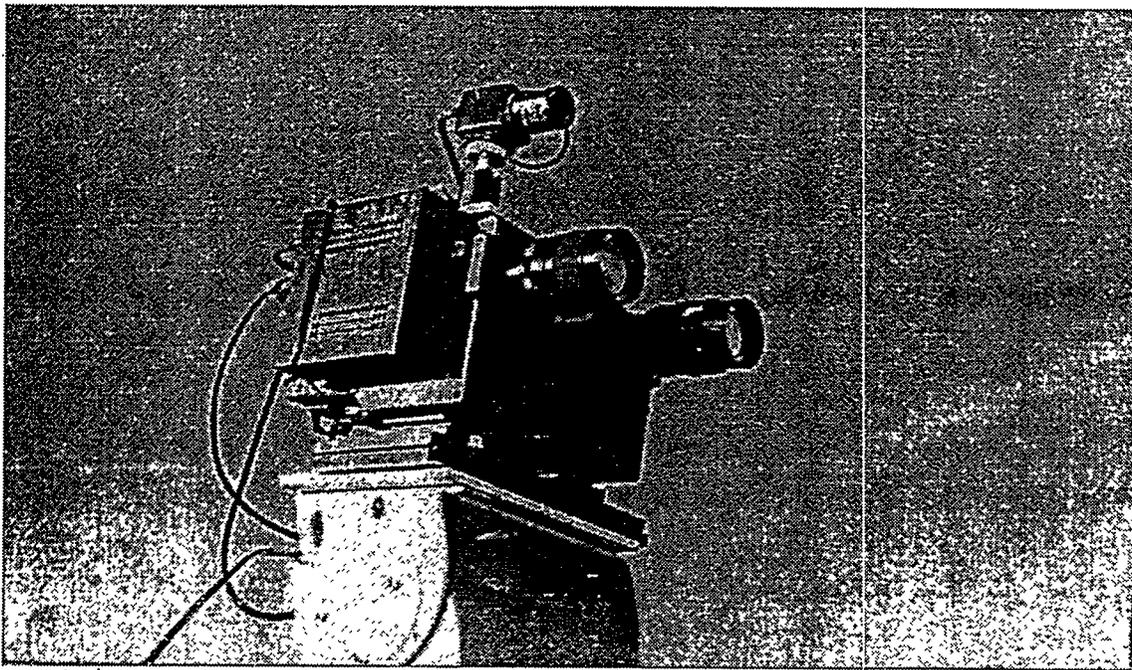


Figure 9. Infrared and visible cameras

which moved along roads between the rows of mines. The video camera used for visible imaging is also shown. The IR cameras were part of the data collection and data pre-processing system shown in figure 10.

This system was located in a motor home connected to the lift truck by an umbilical cable bundle as it moved through the mine field. The Agema computer was used for image acquisition and quick-look and the VAX 4000 Workstation was used for image registration. A PC 486 clone was used for weather data acquisition and surveyed mine ground truth data acquisition, The Macintosh IIFX was used for visual image acquisition, general purpose document preparations and spreadsheet operations. Figure 10 shows the layout of the system and the locations of its components, either in

the boom truck or in the motorhome.

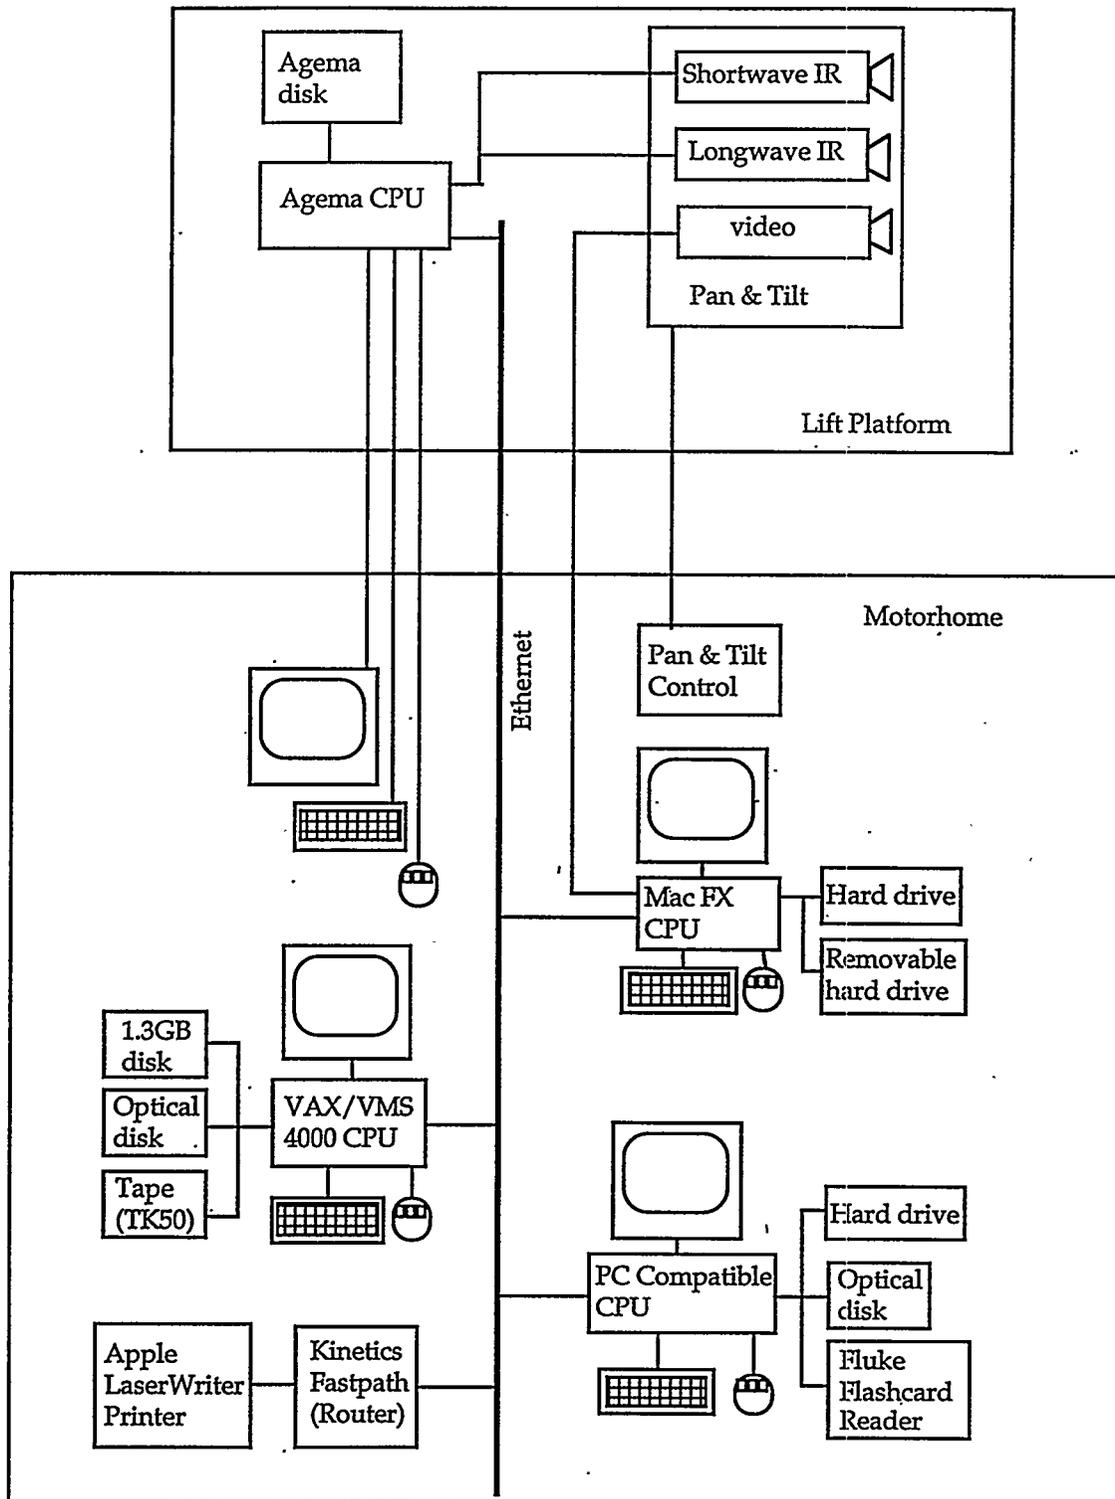


Figure 10. Schematic diagram of infrared and visible camera system

The Agema CPU as well as its secondary storage device (either a magnetic or optical hard disk) were in the boom truck along with the sensors. An "umbilical cord" which consists of the Agema keyboard, mouse, and monitor cables, a thinwire Ethernet cable, the video camera cable, and the pan & tilt control cable connected the boom truck with the mobile home. Inside the mobile thinwire Ethernet connects the VAX, Mac, PC, and Kinetics Fastpath chassis which is the network router. An Apple Laser-Writer Printer was connected to the Fastpath router via PhoneNet.

File services between the Agema and the VAX were provided using Multinet software running the NFS protocol. Sending files between VAX and Mac was done with "PathWorks for Mac" software running the Appletalk network protocol. VAX/PC transfers use "PathWorks for DOS" software running the LAST protocol. Terminal emulation among machines was done with Telnet.

UV, Visible and Near-Infrared and Hyperspectral Sensors

Physical principles

The UV, visible, and near IR cameras and Hyperspectral Imager all image the light reflected from surfaces. Differences in the overall reflectance relate to image brightness, while differences in the spectral dependence of reflectance relate to "color". The spectral reflectance is typically different for different objects. This feature can be put to advantage for enhancing contrast between surface mines, the natural background and other objects or vegetative cover. Daytime visible and near-IR imagery can help to find sites which represent surface disturbances or objects. When these sites coincide with the positions of thermal clutter, detected independently and coregistered with images taken using a thermal IR camera, they provide a valuable aid in identifying, and removing images of objects that otherwise might be confused with a buried mine site on thermal DBIR imagery.

Instrumentation

The UV/Visible/Near-Infrared Sensor Instrumentation data collections were made by Science Applications International Corporation (SAIC) with their three camera system which had previously been used for investigating the detection of surface mines in the DARPA mine/Countermining Program and at the buried mine detection test site at LLNL. For these measurements the 104mm lens on the UV camera was replaced by a 50mm lens borrowed from CSS to accommodate a larger field of view. This provided a field of view of 10 to 15 feet with a resolution of 1/4 inch with a camera elevation of 40 feet. UV, Visible, and Near IR images were recorded for mine field plots H,I,J,K, south and L,M,N, south, and T,U,V,(search plots) south. The images were taken in NTSC standard video and recorded on video tape with individual images for processing provided by a frame grabber. In addition to the buried mine data imagery was obtained for mines placed on the ground surface. There was an apparent advantage of the UV camera in detecting the surface mines at glancing angles and as previously seen, the UV images also showed a reduced shadow. Details of the camera system, measurements and analysis done by SAIC are described in the enclosed SAIC Report: UV Mine Detection Field Test Data.

SAIC also made the hyperspectral data collections with their Hyperspectral Imager (HSI). Data acquisition consisted of recording the spectral content of spatial information averaged over a particular spectral interval and spatial dimension of the object. The resulting data set consists of a spectrum for each pixel of the image. Because of its three dimensional character (two spatial and one spectral), such a data set is often referred to as an image cube. For these measurements the spectral range of 400nm to 930nm was divided into 64 spectral bins with a width of about 8nm. The instrument and results are described in detail in the enclosed report SAIC: Hyperspectral Imaging of Emplaced Land Mines. Image cubes of background plots and buried and surface mines were recorded and are archived but because of the enormous volumes of data produced by this method, it was only done for a small subset of the minefield.

Both sets of the SAIC data are archived at LLNL and is available to interested investigators.

Ground Penetrating Radar Sensors

Physical principles

Data fusion of both dual band infrared (DBIR) and ground penetrating radar (GPR) is anticipated to provide greater detection and discrimination probabilities. GPR has the role of metallic mine detection primarily but may also aid in the detection of non-metallic, i.e. plastic, mines.

Radar systems generally rely on the reflection of long wavelength (> mm) electromagnetic waves by discontinuities in conductance and dielectric constant of the transmitting medium. Although far less transmissive than air, the ground can be used as a transmitting medium. Discontinuities in the ground, should, depending on their nature, provide reflections detectable by a radar system. Such "ground penetrating radar" systems are often close coupled to the ground in order to avoid the problems and losses which result from reflection at the air-ground interface. However, Seasat images showing subsurface structures demonstrated at least the possibility of using above ground imaging radar systems to see subsurface objects.

Usually, mines represent near surface discontinuities. The observability and detectability of mines with radar depends both on the nature of the mine and the background of other non-mine inhomogeneities.

SRI GPR instrumentation

SRI International acquired GPR data for all 34 mine plots at vertical and horizontal polarization with their ground-based ultrawideband (UWB) synthetic aperture radar (SAR). Data were also taken on other plots containing test objects and calibration spheres. All of the metal mines were detected and easily identified. [32] The SRI GPR system, data processing and results are described in the enclosed SRI report: Ground-based Ultrawideband Synthetic Aperature Radar Measurements. The SRI data are archived at LLNL and are available to interested investigators.

LLNL GPR instrumentation

LLNL developed a wideband GPR system similar to the one used by SRI Interna-

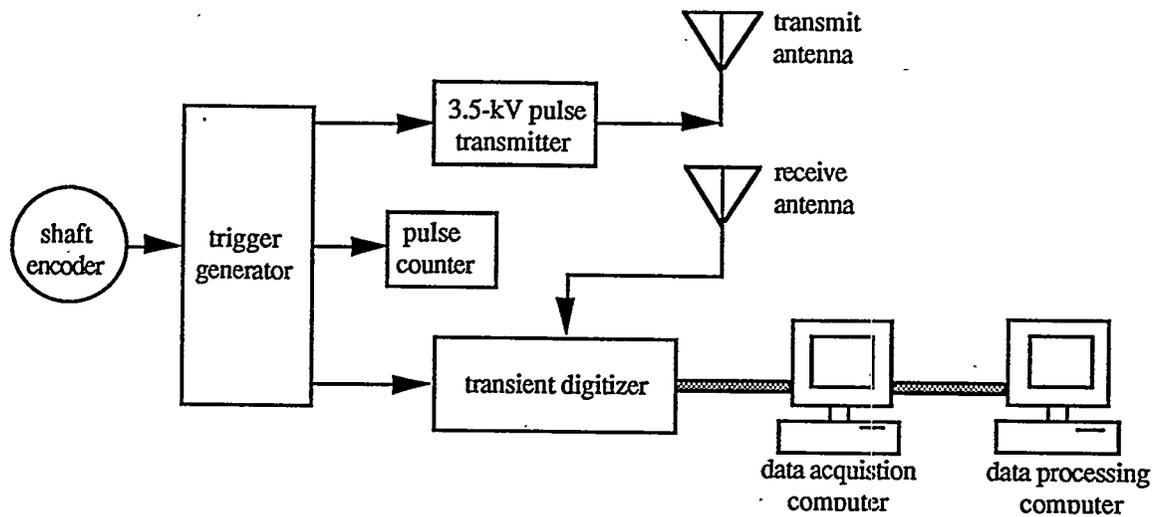


Figure 11. Block Diagram of LLNL GPR system

tional. Figure 11 shows a block diagram of the LLNL GPR system. The current LLNL GPR system combines the transmitter, antennas, and data acquisition and data processing systems in a single van. A HP54720D digitizer is used for recording and a Sun SPARC LX for SAR processing in the field.

A high voltage pulse generator is used for transmitter excitation and a high speed digital data acquisition system for the receiver. Two separate antennas are used for the transmit and receive functions. The transmit antenna is a commercial double-ridged horn and the receive antenna is a custom-built monopole corner reflector that uses a metallic folded triangle element. The receiver system is configured to acquire the scattered energy in the 250MHz to 950MHz frequency range. 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 area. The antenna system was mounted 14 ft to 30 ft above the ground with a depression angle of 30 degrees to the center of the plot. Digitized data from the acquisition unit is transferred to the computer for post-processing and generation of the radar image. Image processing utilizes conventional SAR techniques.

Metal mines buried at depths up to 2" were detected with a signal to clutter ratio between 4:1 and 6:1 and large metal disks were detected at depths ranging from 5 cm to 40cm. The LLNL GPR system, field measurements, and results are described in more detail in the Appendix A.

Other Data Collections

In addition to measurements sponsored by this program, measurements have been made by other investigators. MIT/Lincoln Laboratory (sponsored by The Army Research Laboratory) made measurements on a limited number of mine plots with their stepped CW rail-SAR GPR system. ERIM made measurements of infrared spectra from soil above buried mines and background regions with an infrared spectrometer.

Neither of these data sets have been made available to us at the time of this writing. More information on these data can be obtained from Denis Blejer at MIT/LL (617)-981-3455 and John Selden 313-994-1200 x 2361 at ERIM.

IV. Data Processing

LLNL Data Acquisition and Pre-Processing

Computational data analysis is important to standoff mine detection efforts for two reasons. First, in order to be effective, a mine detection system must be able to cover a large swath of ground and quickly and reliably report on the presence of mine fields. An automatic target recognition system is desirable because it will increase the rate at which images can be tested, it does not suffer from fatigue or boredom, and it does not rely on the subjectivity and training of a particular operator. Thus the mechanization of the process improves its efficiency. The second reason is that it is very difficult for humans to visualize multisensor data quickly and effectively. Image processing and multisensor fusion techniques allow for the efficient compression of a number of distinct images of a scene into a "superimage" which provides much improved signal to noise ratios for the desired targets than single sensor images, improving the effectiveness of both computer and human operators.

Image 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 polynomial warping algorithm, which performs translation, rotation, scaling and perspective corrections to the images. 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. The following section describes the flow of LLNL data for the Dual Band Infrared (DBIR), visible, and Ground Penetrating Radar (GPR) measurements

Dual Band Infrared:

Figure 12 shows the flow of DBIR and visible data from capture to registration. The DBIR data are captured with the Agema infrared system. Short Wave Infrared data (SWIR) of 3-5 microns and Long Wave Infrared (LWIR) of 8-12 microns are read from a pair of cameras mounted 60 feet above the ground in the center of the roadway about 25' from the center of the subplot. A single 18' by 18' subplot is viewed at a time. The raw SWIR and LWIR images are then moved to a VAX workstation for registration where the dual band data are registered to "plan view", that is, a view directly overhead above the center of the subplot. A database of surveyed reference positions for the given subplot are read. These reference positions, pointed at by the arrows in the Raw images, are used with polynomial warping and/or perspective warping algorithms to map the the raw image to the "plan view". The registered images are then prepared for distribution and archived.

Visible data:

The visible wavelength data are captured using either a video camera connected to a Macintosh equipped with a frame grabber and custom acquisition and control software or digitized photographic images. The captured images are then moved to the VAX workstation and registered to plan view in a manner similar to the DBIR. These images are also archived and networked to the ATR team.

GPR data:

Raw GPR images are acquired by a custom system built at LLNL. The images are built as a series of voltage versus time radar echo return signals captured along successive positions of a synthetic aperture. Figure 13 shows the flow of the GPR data from acquisition to target recognition. The raw image (only a portion of which is shown in the figure) is preconditioned using ground range correction, matched filter integral, and/or time domain jitter correction. The preconditioned images are then "focussed".

Objects echoing along the synthetic aperture report as hyperbolae in the raw image. Given parameters such as transmitter/receiver positions with respect to the sub-plot being viewed, synthetic aperture distances, and the frequency range of the antenna, the focussing operation integrates the hyperbolae to reconstruct the image along the ground plane. The focussed images are then rescaled (magnified) according to ground truth values, shifted, and rotated (i.e. registered) to plan view. The registered GPR images are then archived and processed further to detect and identify targets and fused with the DBIR images.

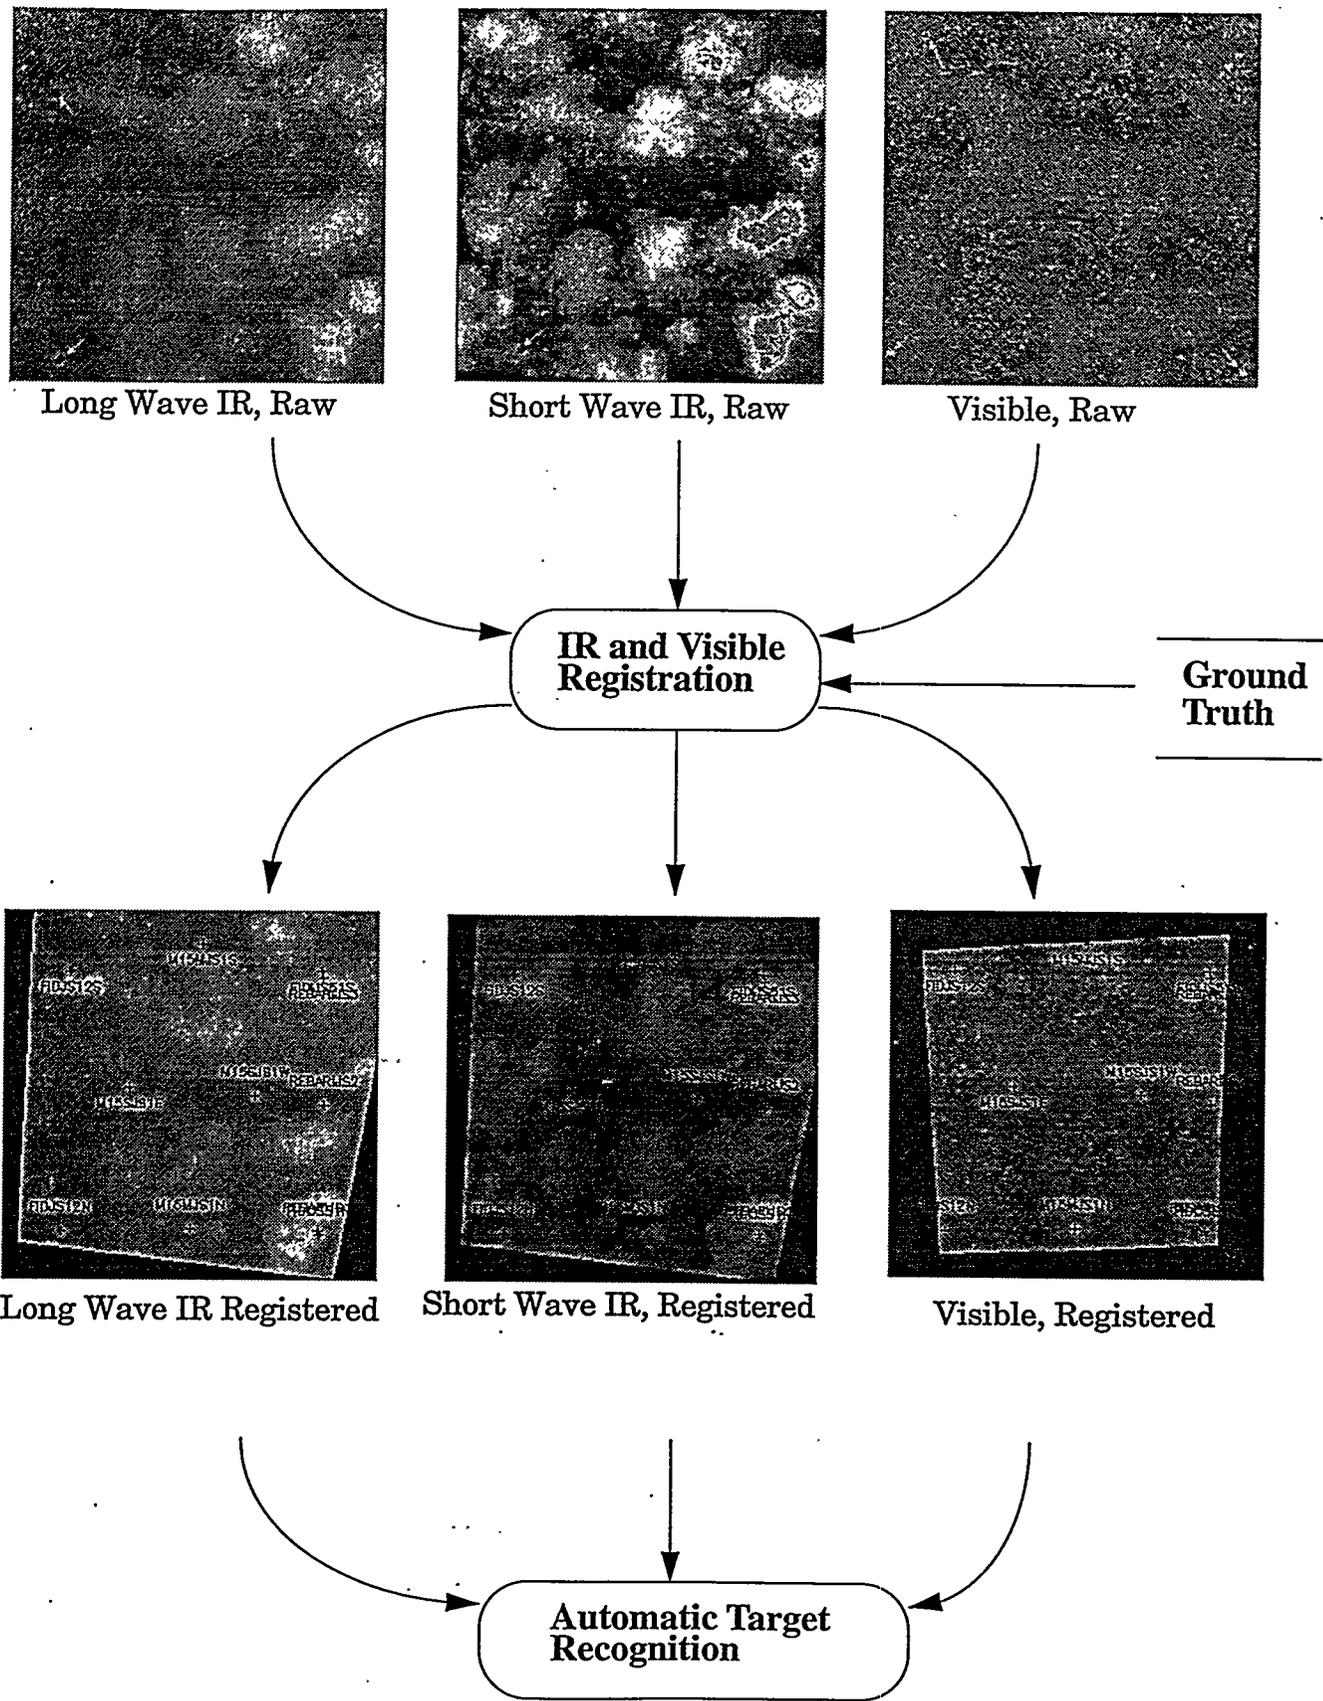


Figure 12. Infrared and visible data flow from capture to ATR processing

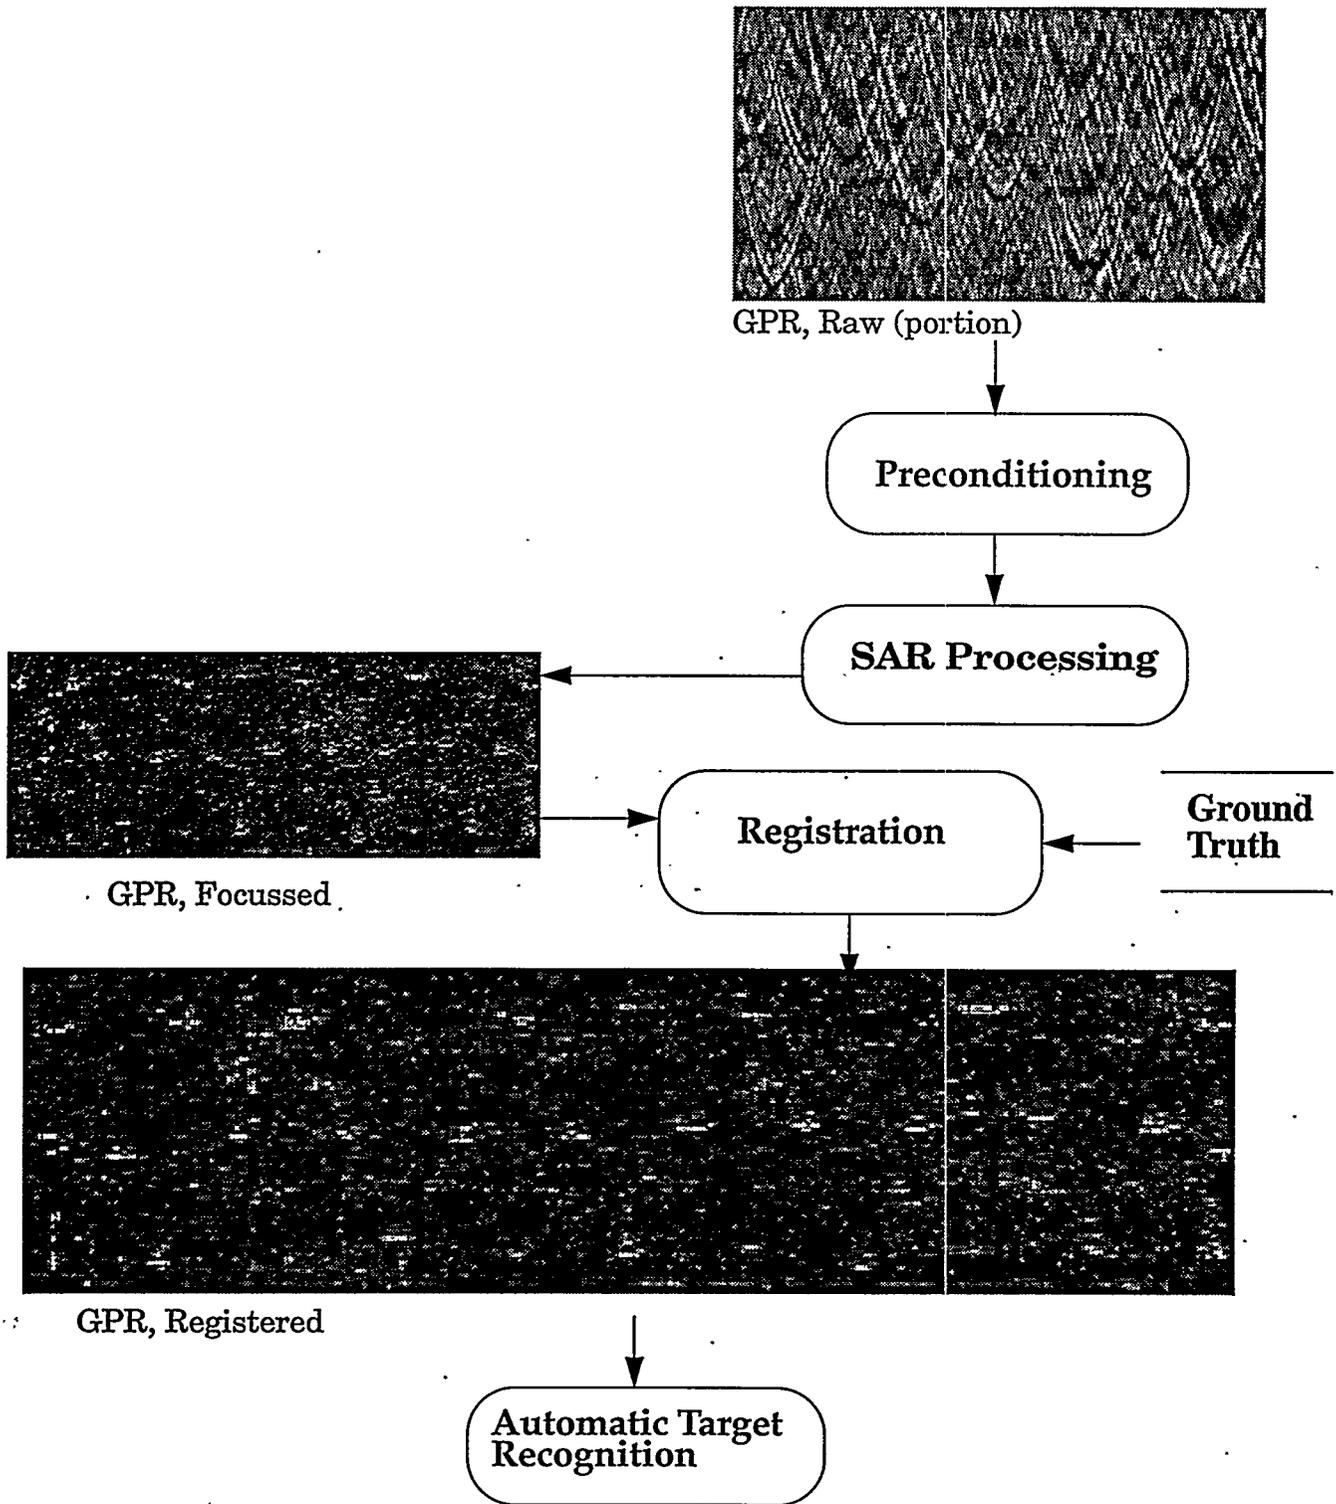


Figure 13. Flow of GPR data from acquisition to ATR processing

Automatic Target Recognition and Data Fusion

Past research has shown that it is extremely difficult to distinguish buried mines from background clutter in images obtained from any single sensor. We believe that data fusion applied to simultaneous observations with a variety of different sensors may provide a reliable signature for buried mines. This work focuses on the fusion of images from two infrared sensors with different band passes. We use feature-level fusion and supervised learning with the probabilistic neural network (PNN) to evaluate detection performance. The novelty of the work lies in the application of advanced target recognition algorithms, the fusion of dual-band infrared images and evaluation of the techniques using two real data sets.

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]. We have also implemented the algorithms in IDL, using functions written in C.

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. Examples of pseudo-images are temperature and emissivity images calculated from IR intensity 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 algo-

rithm 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 ratio" (E) and "temperature ratio" (T):

$$E = \ln\left(\frac{S/s}{L/l}\right)$$

$$T = \ln\left(\frac{(L/l)^2}{S/s}\right)$$

where S is the short-wavelength intensity, s is the average value of the pixels in S, L is the long-wavelength intensity and l is the average value of the pixels in L. The emissivity ratio should be insensitive to the differences in temperature of a single substance, but should show emissivity contrasts among different substances. The temperature ratio should isolate differences in temperature of a substance and be less sensitive to differences of emissivity.

Feature Extraction

After the pre-processing has been performed, the sub-image samples (or tiles) in the image are selected. These sub-images are square regions (N x 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].

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 or irrelevant 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].

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 situation specific empirical knowledge 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 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 [6], and the probabilistic neural network [8, 16].

Supervised Learning Results for 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. We conducted controlled experiments at

the LLNL site. The soil was prepared so the surface was cleared and free of vegetation, stones and other clutter sources. Additionally, the data were acquired in both the daytime and nighttime..

In this experiment, we defined a two-class problem. The first class (called "Mines") corresponds to buried mines and filled holes. (A filled hole is one for which the soil was removed and then replaced.) Filled 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 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 those of mines [29-30].

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. These 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 cases 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.

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 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 x two images), there were 4095 (212-1) possible feature subsets tested for each case. The feature subset which provided the highest probability of detection was selected as the "optimal" subset. The image labeling along with the post-processing are described in the next section.

Image Labeling and Post-Processing

After the feature selection had been performed, the "optimal" subset of features was computed for every pixel in the testing images. These feature vectors were then classified as belonging to either the class "Mine" or "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 labeled as 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.

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 2.1 and 2.2.

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, 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. 2.1. Feature selection and image detection results using only the long and short wave-length 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. 2.2. Feature selection and image detection results using only the temperatures and emissivity images

Supervised Learning Results for a Cluttered Environment

The data for this experiment were acquired from the NTS site. This site, which is 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 and disturbance of the soil surface was minimized.

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 bushes could be determined, and eliminated from further consideration as a mine by forming a mask. 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 ground similar to where mines were actually buried. The square 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. The feature vector for 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 the previous section, which analyzed the long and short wavelength images separately from the temperature and emissivity images.

Only nighttime images which contained metal mines and surrogates have been processed. These results are presented below.

Feature Selection

The feature selection was accomplished using 1480 background samples (sub-images), 60 mine samples, and 45 surrogate samples collected from the database of images taken over the test plot. 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 estimated probability density functions of the three features

shown in Fig. 14.1. The image labeling using these features is described in the next section.

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 (PNN) [8]. The output was an image in which 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 14.2 to 14.4. Figure 14.2 displays the long wavelength, short wavelength, temperature, and emissivity images for a plot. Figure 14.3 displays the probability image which results after the classification. The posterior probability of the class "mine" is plotted at each pixel. The brighter areas of this image indicate locations for which a high probability of a mine exists. Lastly, Fig. 14.4 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.

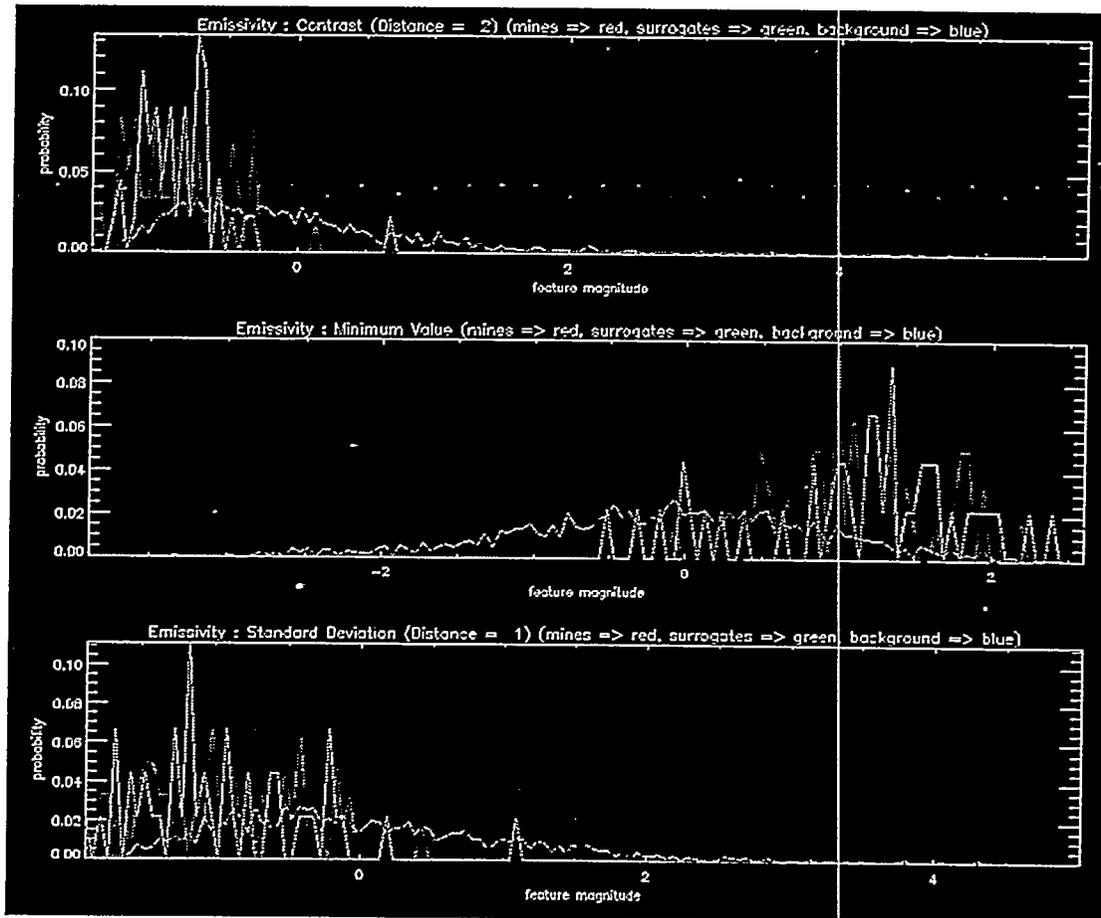
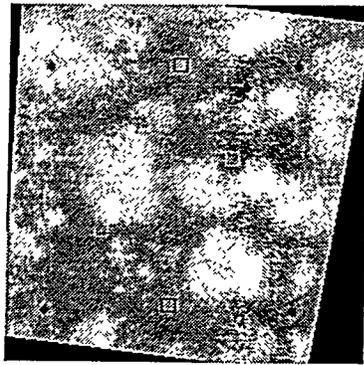
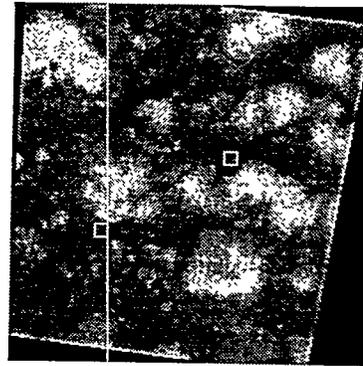


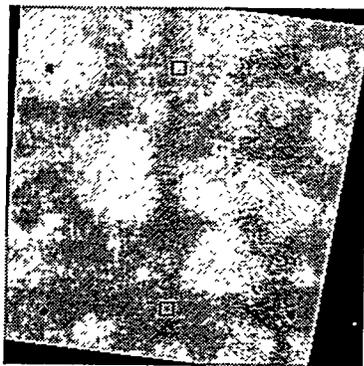
Figure 14.1. Probability density function estimates (histograms) of the three "best" features. The smoothest curves represent the background statistics. The others represent mine and mine surrogate statistics.



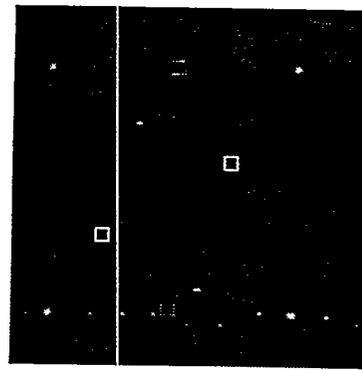
(a) Long wavelength image



(b) Short wavelength image



(c) Emissivity image



(d) Temperature image

Figure 14.2. Long wavelength (a), short wavelength (b), temperature (c), and emissivity (d) images. Locations of mines and mine surrogates are indicated by squares.

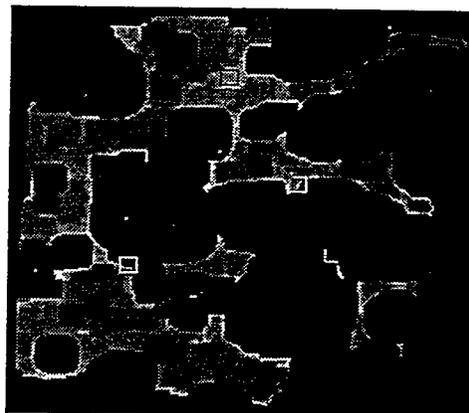


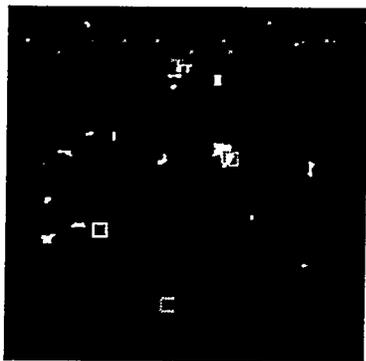
Figure 14.3. Image representing the probability of a mine or surrogate pixel. The black areas were masked out as they were identified as vegetation.

The results for the cluttered environment at the NTS were much less favorable than the results for the controlled experiments at LLNL (uncluttered). At the NTS, the IR images of mines were relatively weak compared to clutter, so only limited results were obtained

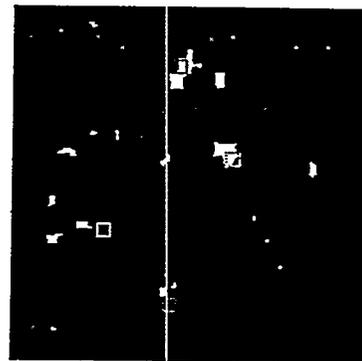
Discussion

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 back propagation neural nets. Using PNN we could distinguish between holes with and without mines.

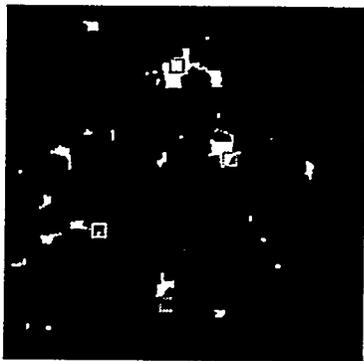
However, further work with cluttered fields provided considerable instruction as to the utility of data fusion techniques. Current results indicate that the best performance was achieved using texture features extracted from the pseudo-image "emissivity". We currently believe this to be physically related to the surface disturbance resulting from the mine emplacement. It is clear that even with these results, additional sensors must be included for reasonable detection reliability to be achieved.



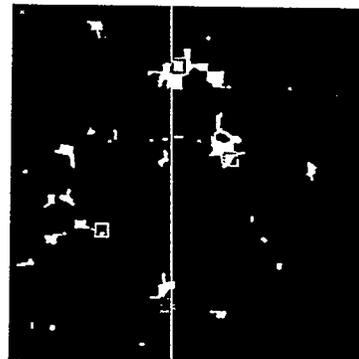
(a) most likely 500 pixels



(b) most likely 1000 pixels



(c) most likely 1500 pixels



(d) most likely 2000 pixels

Figure 14.4. 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.

V. Summary of Results

In the course of data collection at NTS more than 500 Mb of IR, GPR, and other electro-optic data were collected over a variety of conditions. A complete analysis of all this data is beyond the scope of this project, and only part of the data (specifically the IR data) were used as input for the complete automatic target recognition sequence. On the other hand, we have surveyed the data to provide a basis for general conclusions. This section summarizes these general conclusions resulting from the data collections and analysis of this project.

Visual inspection of data

For basic 2-D image analysis, visual inspection is still an extremely powerful technique. In spite of its drawbacks, it allows a rather high speed analysis of large volumes of data with the fairly sophisticated pattern recognition and data fusion capabilities of the human being.

The figures referenced below are marked to show the locations of mines and fiducial markers. The labels carry information on the mine type.

Visible images

A walking tour of the Mine Test Facility demonstrates that buried mine detection is a difficult job, especially mines in a weathered environment. Even with the prior knowledge of the approximate locations of the mines, it is quite unlikely that one will locate more than a few. Figure [6] shows how little visible signature is present to the eye. However, as shown in Figure 15, there is evidence of subtle color changes due to

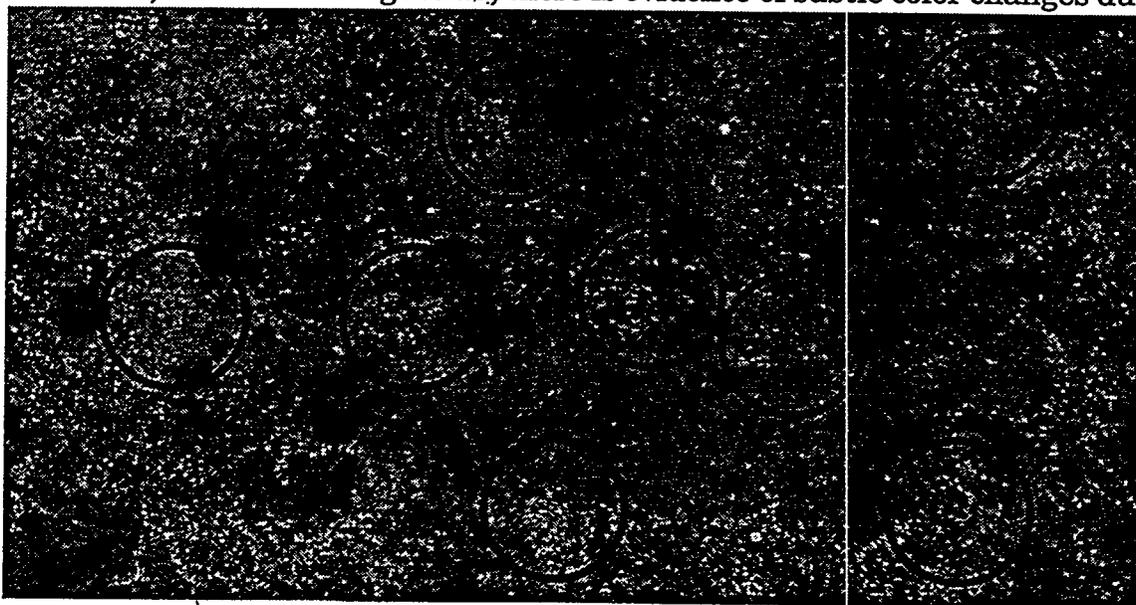


Figure 15. Visible image of sub-plot MS2

surface disturbance by human activity in the area and the mine emplacement hole which could be detected by hyperspectral sensors. Fusion of multiple sensors would aid in reducing the false alarm rate.

While the initial investigations of DelGrande, et. al. [11] showed that buried mines have detectable thermal signatures, these signatures are easily masked by the sizable temperature variations due to natural clutter. Figures 16 and 17 show several views of mine plots in the SWIR and the LWIR. These, images, taken at night, show that bushes and vegetation, which tend to equilibrate to the air temperature, are significantly warmer than the ground. The temperature variations in the image due to these plants is significantly larger than the variations due to the mines.

The mine signatures, in fact, are generally slight depressions in the temperature. These depressions are not always present, but they can be seen in several of the mines marked in JS2 and JS3 images. What is clearly the case is that mines will not be found simply by setting thresholds on temperature.

GPR

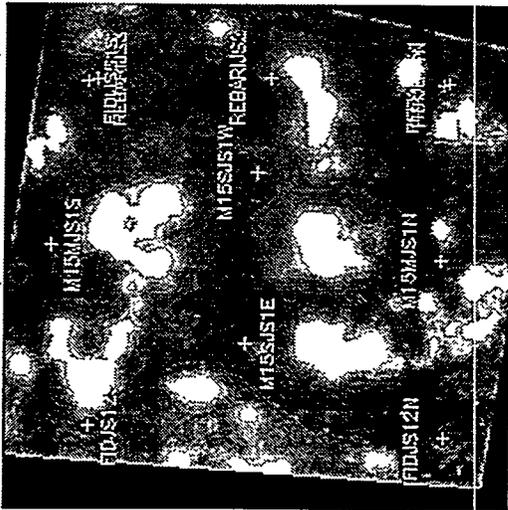
LLNL GPR results are described in detail in Appendix A. The GPR images typically show metal mines clearly. Large metal test objects were imaged with high signal-to noise (SNR) and signal-to-clutter ratio (SCR) to depths of at least 20". Large plastic objects can be imaged with a lower SNR and SCR. Medium plastic objects are difficult to image. All of the metal mines buried to depths of 2" were easily detected and easily identified. The reconstructed SAR image of plot I South in Figure 18a shows that metal mines and surrogates are clearly detected. This result is important as it is the first case where standoff radar systems have been shown to detect mines in a natural, cluttered environment. Results from Project Ostrich [31] showed mines clearly only when they were on the surface. Also rising above the background are the metal fiducials, which were chosen to be visible to the radar system.

Images with plastic mines show that detection using reconstruction based only on reflected power will be difficult. Although some plastic mines or surrogates are visible in Figure 18b showing the reconstructed SAR image of plot K North, the signal to clutter ratio is low.

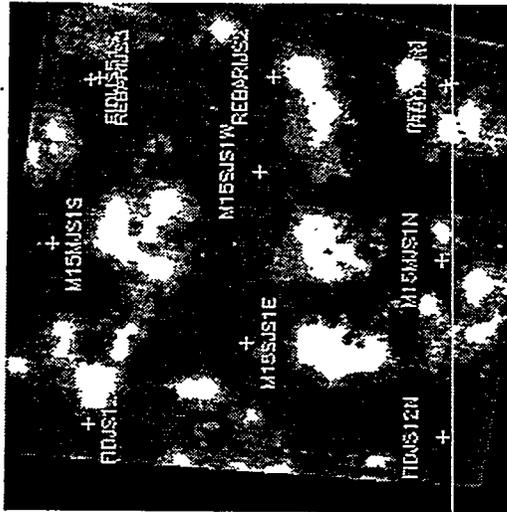
SRI International GPR results[32, Enclosure 3] were similar to those above. All of the metal mines and mine surrogates were detected with high signal-to-noise and signal to clutter ratios. All of the plastic surrogates and some of the large plastic mines were also detected. SRI also acquired data for two different sizes of dc calibration spheres suspended in the air. This allowed processed SAR images to be calibrated in absolute radar cross section (RCS) in dBsm.



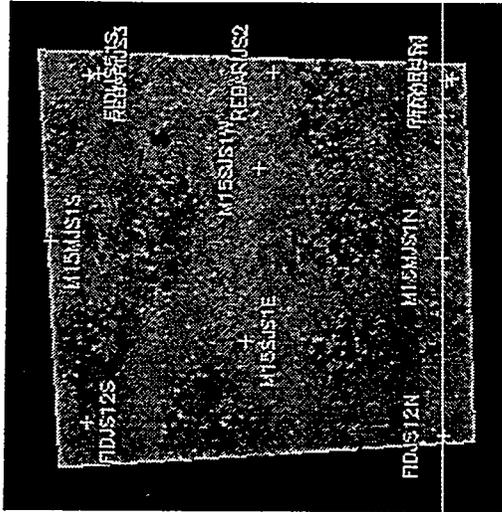
Long wave IR composite of J south



Long wave IR

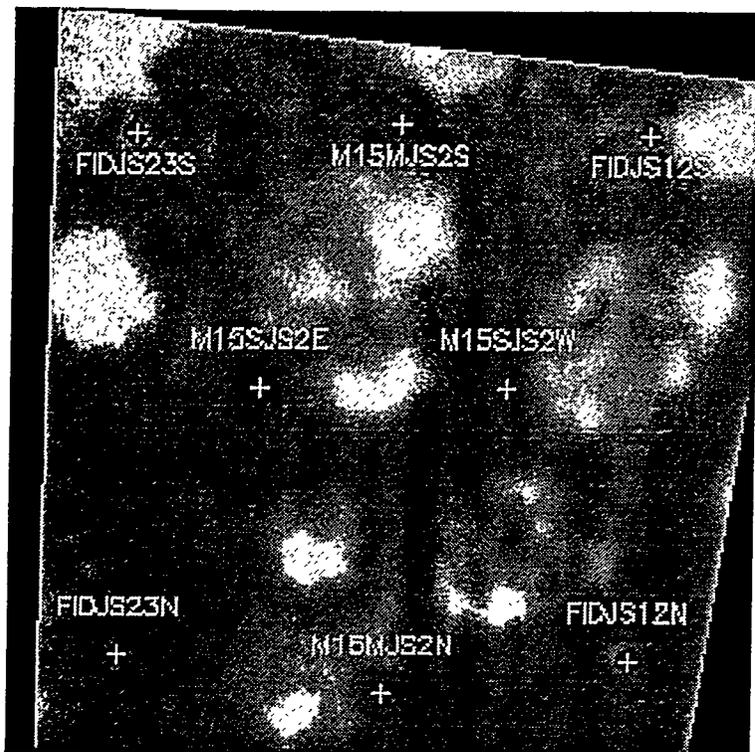


Short wave IR

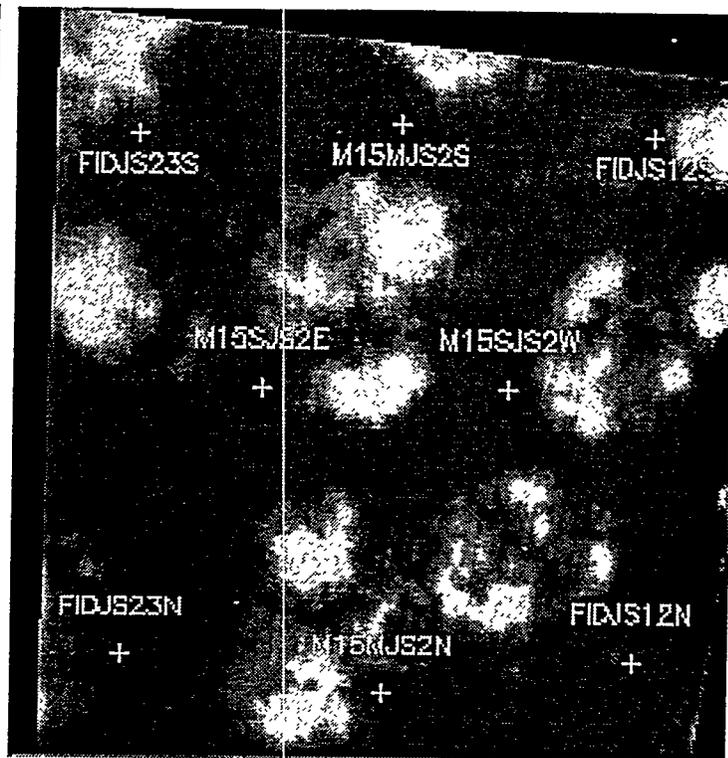


Visible

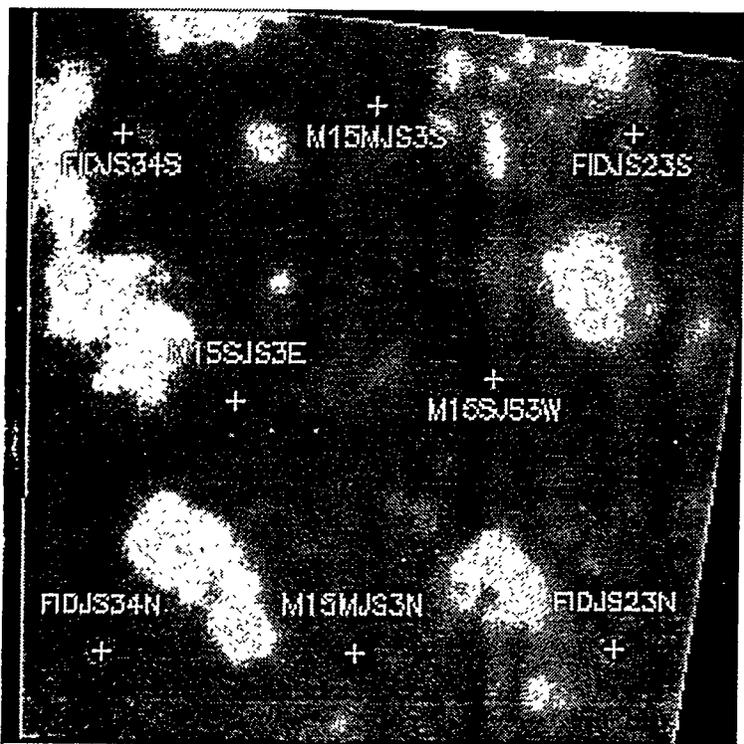
Figure 16. Infrared and visible images of plot J South



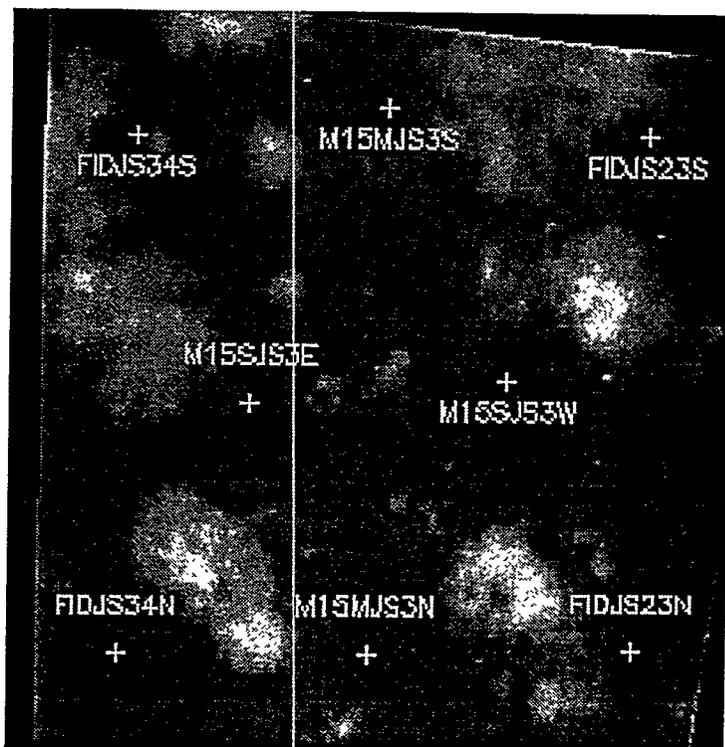
JS2 LWIR



JS2 SWIR



JS3 LWIR



JS3 SWIR

Figure 17. Infrared Images of Subplots JS2 and JS3

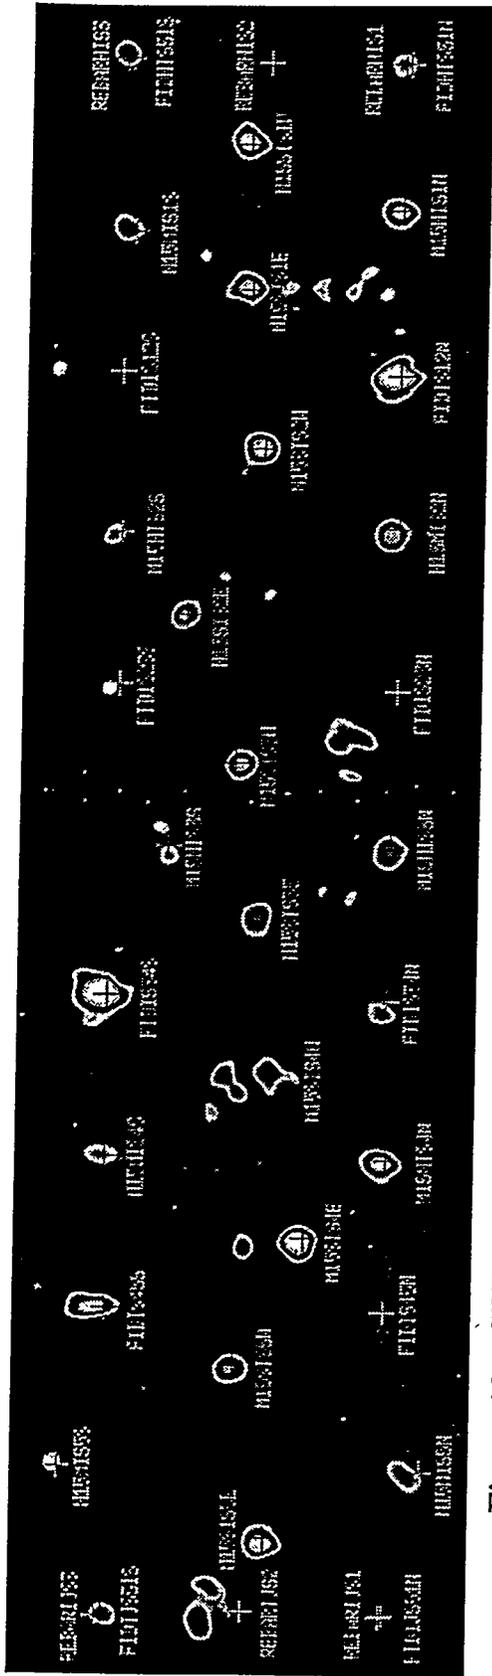


Figure 18a. GPR image of plot I south showing M15 metal mines and surrogates

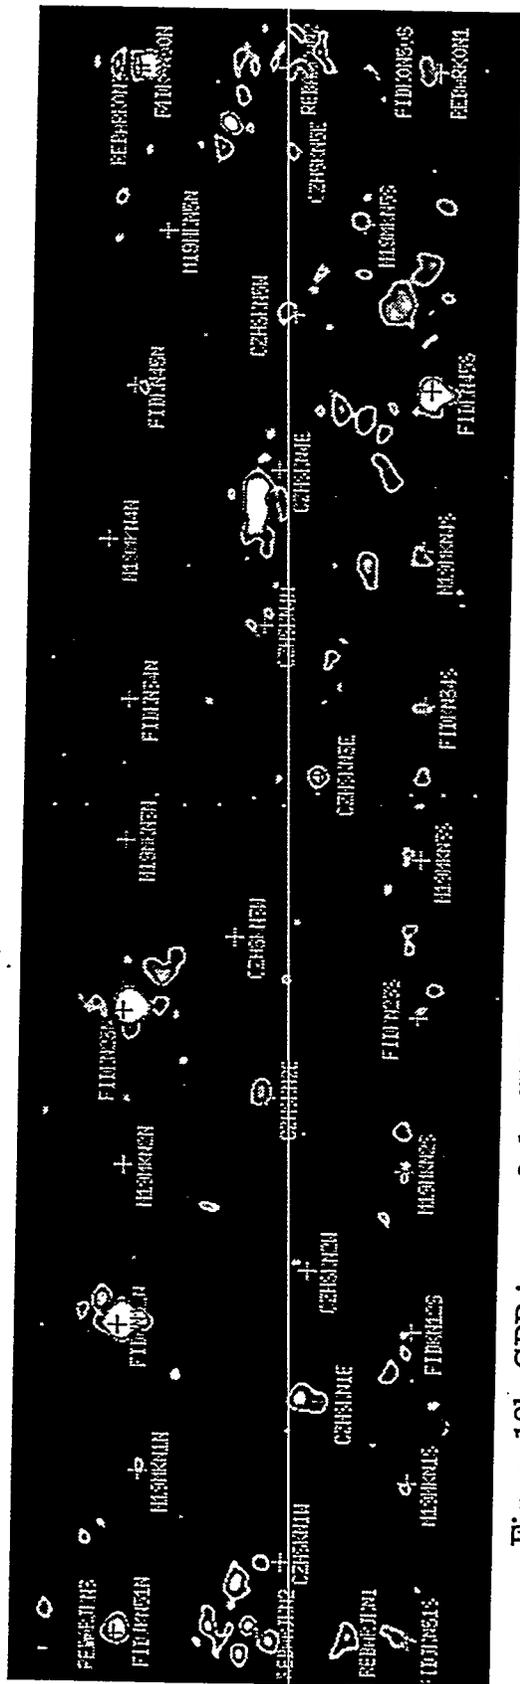


Figure 18b. GPR image of plot K North containing M19 plastic mines and Czech surrogates

Summary of ATR results

The ATR analysis for the controlled experiments at the LLNL site gave excellent detection results with good false alarm rates. However the detection and false alarm results from the natural NTS environment were not adequate for robust detection. The second IR band improved detection performance but the primary problem remains the inability to discriminate between relatively weak mine signatures, when present, and the rather high level of natural clutter. As discussed below the ATR analysis for COBRA data was more successful

The ATR analysis discussed in the previous section supports the conclusion that adding a second sensor channel improves detection performance. In that case, the sensor channel was an additional broad band IR channel. The addition of more sensor channels, either by using multi-spectral or hyper-spectral sensors is one possible avenue for improvement. The fusion of IR data and data from the multiband COBRA sensors discussed below could demonstrate the benefits of this approach.

The addition of detector channels other than IR is also a worthwhile approach, and, as data to support such investigations is available from this program, it may also be a workable one. The fusion of GPR and IR images is a particularly interesting option which should be pursued in the future.

COBRA data

We have applied our methods to data other than that gathered at NTS. In support of the COBRA program, we have analyzed the six band visible-near IR imagery of buried and surface mines which existed as of June, 1994. While this data is not of particularly high quality, we have found reasonable detection probabilities are achievable because signatures of the surface mines were sufficiently strong. Improved data should be forthcoming. A more detailed description of the COBRA analysis is in the report included as Appendix B.

VI. 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 N. K. Delgrande, A. Toor, G. A. Clark, R. J. Sherwood, J. E. Hernandez, S. Y. Lu, P. Durbin, D. Nelson, M. Lawrence, M. Spann, and C. Fry, "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 G. 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, "Sensor Feature Fusion for Detecting Buried Objects," SPIE's 1993 International Symposium and Exhibition on Optical Engineering and Photonics in Aerospace and Remote Sensing, Conference on Underground and Obscured Imaging and Detection, Orlando, FLA, April 12-16, 1993.
- 30 Michael R. Buhl, Jose E. Hernandez, Gregory A. Clark, and Sailes K. Sengupta, "Dual-Band Infrared Buried Mine Detection Using a Statistical Pattern Recognition Approach," Lawrence Livermore National Laboratory report UCRL-ID-

114838, August, 1993.

³¹John V. E. Hanser, Judy Ehlen, Timothy D. Evans, Richard A. Hevenor, Project Ostrich-A Feasibility Study: Detecting Buried Mines in Dry Soils Using Synthetic Aperture Radar, US Army Corps of Engineers Topographic Engineering Center Report -TEC-0040

³²David R. Rutt, Roger S. Vickers, Uses of an Ultra-Wideband System for Buried Mine Detection, To Be Presented at Seventh National Conference on High Power Microwave Technology, 31 Oct-4 Nov, 1994.

Appendix A

Ground-penetrating radar for buried mine detection

Paul D. Sargis, F. Dean Lee, E. Stephen Fulkerson, Billy J. McKinley, William D. Aimonetti

Lawrence Livermore National Laboratory, P.O. Box 808, M/S L-54, Livermore, CA 94551

ABSTRACT

LLNL is developing an ultra-wideband, side-looking, ground-penetrating impulse radar system that can be mounted on an airborne platform for the purpose of locating buried mines. The radar system is presently mounted on an 18-meter boom. We have successfully imaged a minefield located at the Nevada Test Site. The minefield consists of real and surrogate mines of various materials and sizes placed in natural vegetation. Some areas have been cleared for non-cluttered studies. A technical description of the system will be presented, describing the wideband antennas, the video pulser, the receiver hardware, and the data acquisition system. The receiver and data acquisition hardware are off-the-shelf-components. The data was processed using LLNL-developed image reconstruction software, and has been registered against the ground truth data. Images showing clearly visible mines, surface reference markers, and ground clutter will be presented.

2. INTRODUCTION

Various types of synthetic-aperture radar (SAR) systems can be applied to the problem of locating buried mines and minefields. High-altitude airborne SAR systems have the advantage of being able to search large areas quickly, but they require high-power transmitters to couple adequate energy into the ground. Ground-coupled systems can penetrate deeply into the soil, but they require travel over the surface of the site being surveyed and only provide data over a narrow swath. Our "standoff" ground-penetrating radar (GPR) system is mounted on an elevated mobile boom, simulating a slow, low-flying airborne platform. We can examine areas up to ten meters wide with a penetration depth that is between that of airborne and ground-coupled systems.

Instead of using a pulse-modulated swept-frequency technique, we chose to implement an impulse approach, because of its relative simplicity and because of our expertise in impulse generation, transient digital recording, and SAR processing. Our approach is based on some of the techniques developed by SRI International.^{1,2}

3. SYSTEM DESCRIPTION

Our system presently requires a caravan of two vehicles. Transmit and receive antennas are mounted on a boom truck having an altitude capability of 18 meters. The boom truck is followed by a motorhome which houses the instrumentation. As pictured in Figure 1, the two vehicles are joined by a bundle of cables. A simple block diagram of our GPR system is presented in

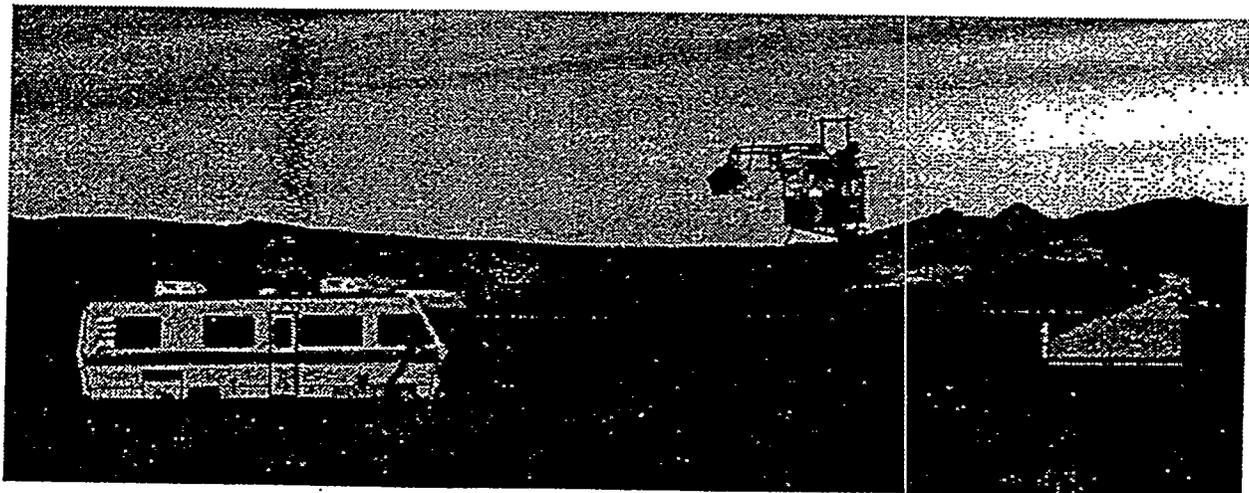


Figure 1. Illustration of GPR vehicle caravan

Figure 2. The trigger source for the system is a shaft encoder mounted on a bicycle wheel which follows the boom truck. A trigger pulse is issued whenever the wheel turns a selected fraction of a revolution. When prompted by the shaft encoder, the trigger generator distributes triggers to the pulse transmitter, a pulse counter, and a transient digitizer with appropriate delays to compensate for cable lengths.

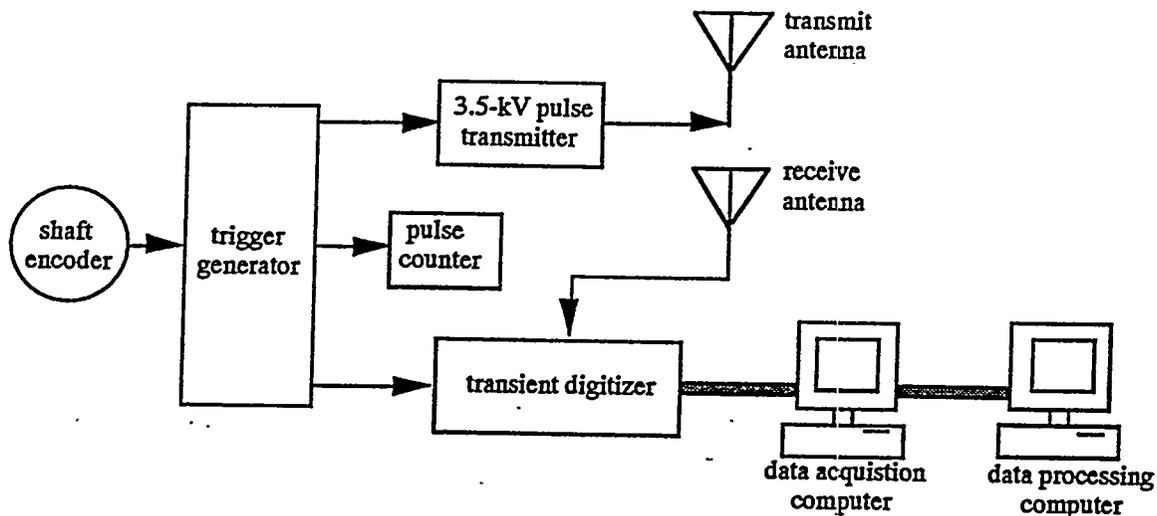


Figure 2. Block diagram of GPR system

The pulse transmitter is a LLNL-designed unit which uses avalanche transistor technology to generate the high-voltage video pulse shown in Figure 3. We used pulse-shaping techniques to optimize the spectral content of the transmitter over the passband of interest. (See Figure 4.) Optimum power transfer occurs when the spectral content of the transmitter matches the response of the transmit antenna. If the transmitter generates significant energy outside the passband of the antenna, it will reflect back toward the transmitter. This energy may then reflect a second time from the transmitter output, causing interfering pulses to be transmitted from the antenna.

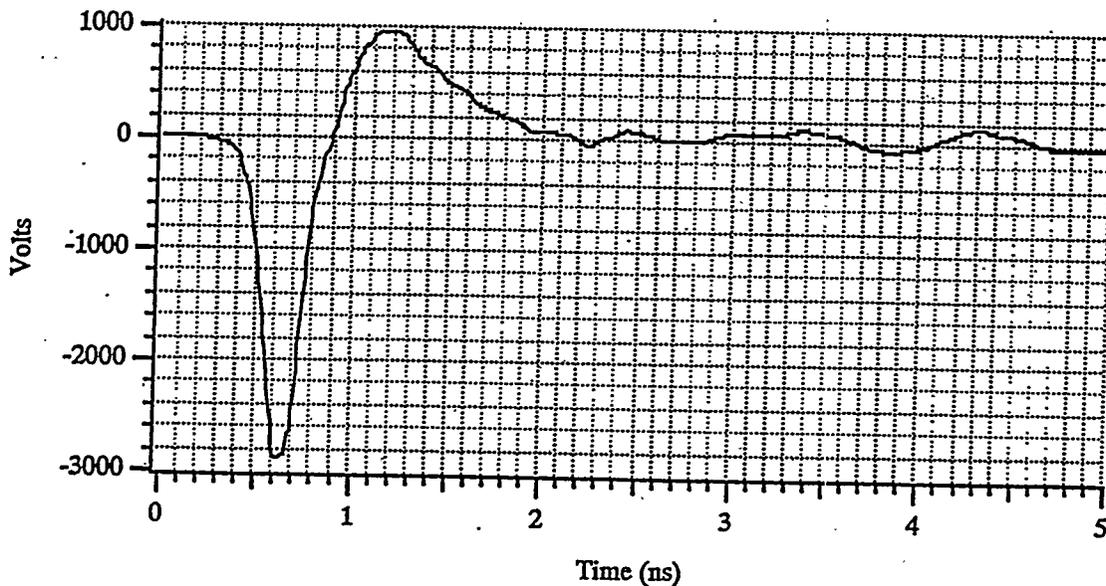


Figure 3. Pulse transmitter output waveform

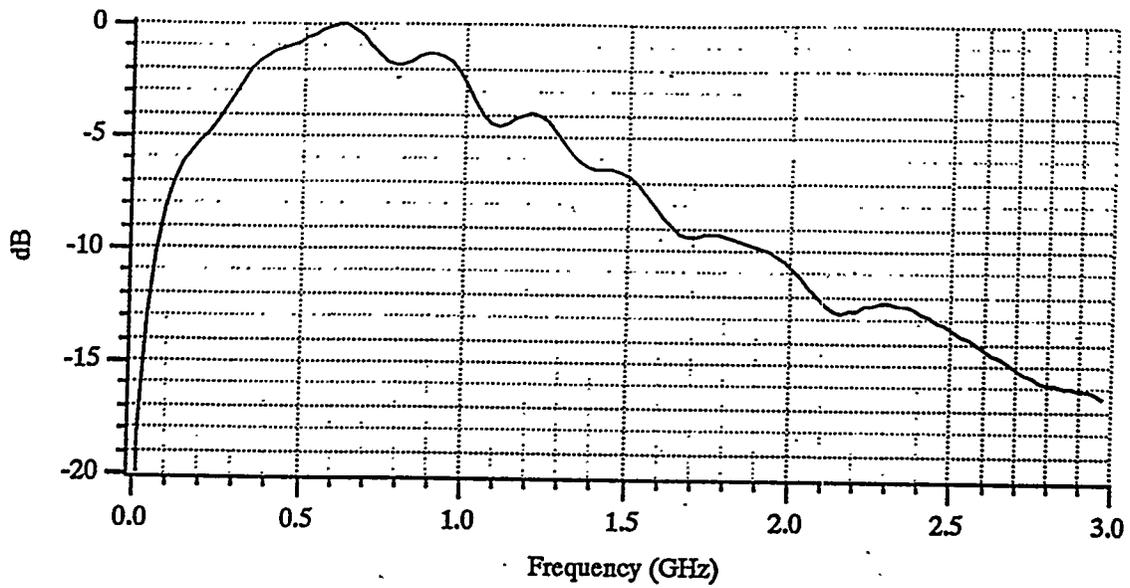


Figure 4. Spectrum of pulse transmitter output

The transmit antenna, pictured in Figure 5, is a commercial double-ridged horn antenna and is designed to operate from 200 MHz to 2000 MHz. Although the 3-dB antenna beamwidth varies with frequency, it is on the order of 45 degrees in both the E-plane and the H-plane.

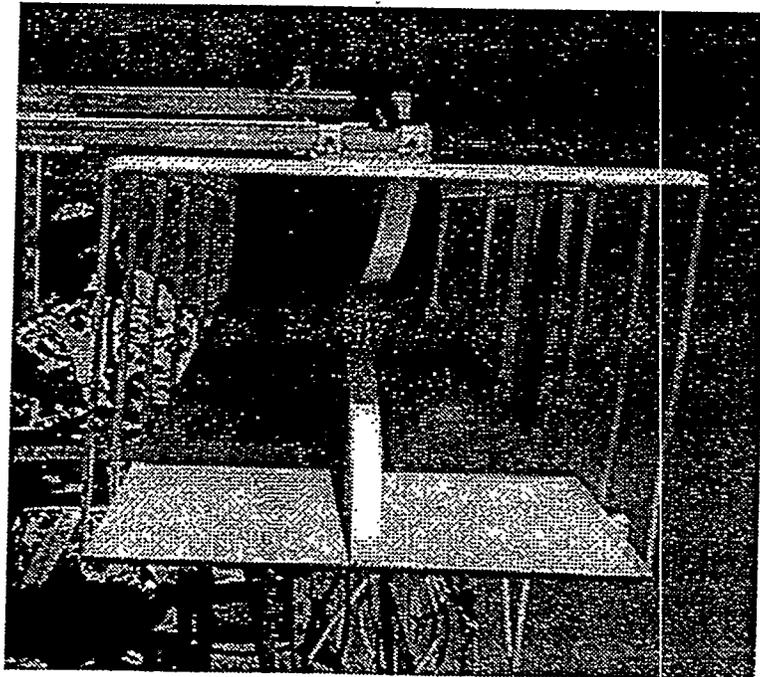


Figure 5. Double-ridged horn transmit antenna

Our receive antenna, shown in Figure 6, is a unique design. The monopole corner reflector uses a folded triangle element to permit operation from 400 MHz to 1500 MHz. This design has several advantages over a conventional resistive dipole corner reflector. It has considerably more gain than a resistive dipole, it is half the size, and it takes a simple coaxial feed with an autotransformer matching network. Its geometry minimizes crosstalk from the adjacent horn antenna. Like the horn antenna, the 3-dB beamwidth for this antenna varies with frequency. It is on the order of 40 degrees in both the E-plane and the H-plane.

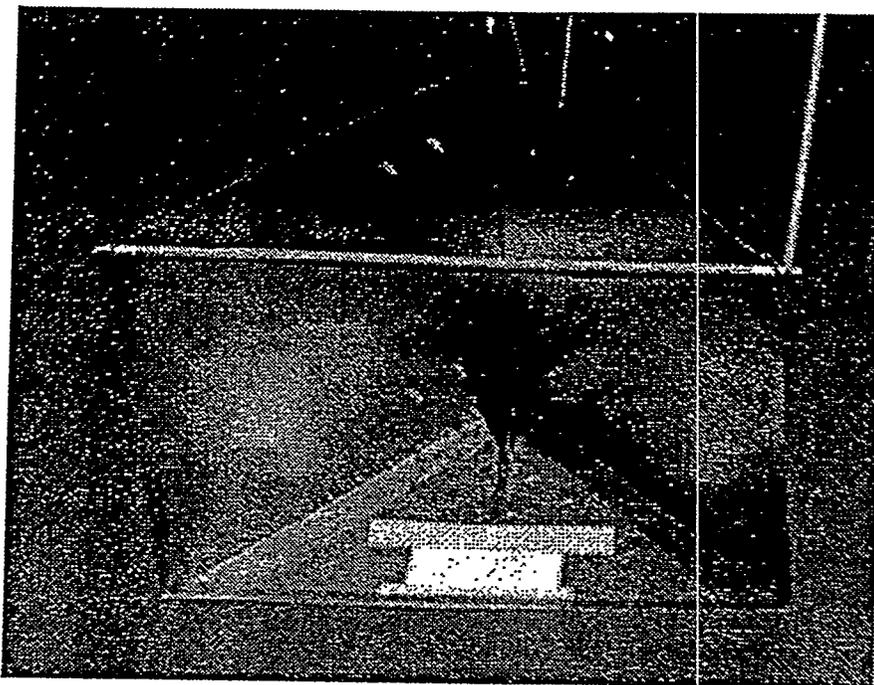


Figure 6. Monopole corner reflector receive antenna

The frequency response of the pair of antennas is plotted in Figure 7. They are capable of efficiently transmitting a pulse having spectral content between 400 MHz and 1500 MHz. Because the monopole corner reflector has a narrower bandwidth than the horn, we use it on the receiver rather than on the transmitter because it rejects local VHF two-way radio signals.

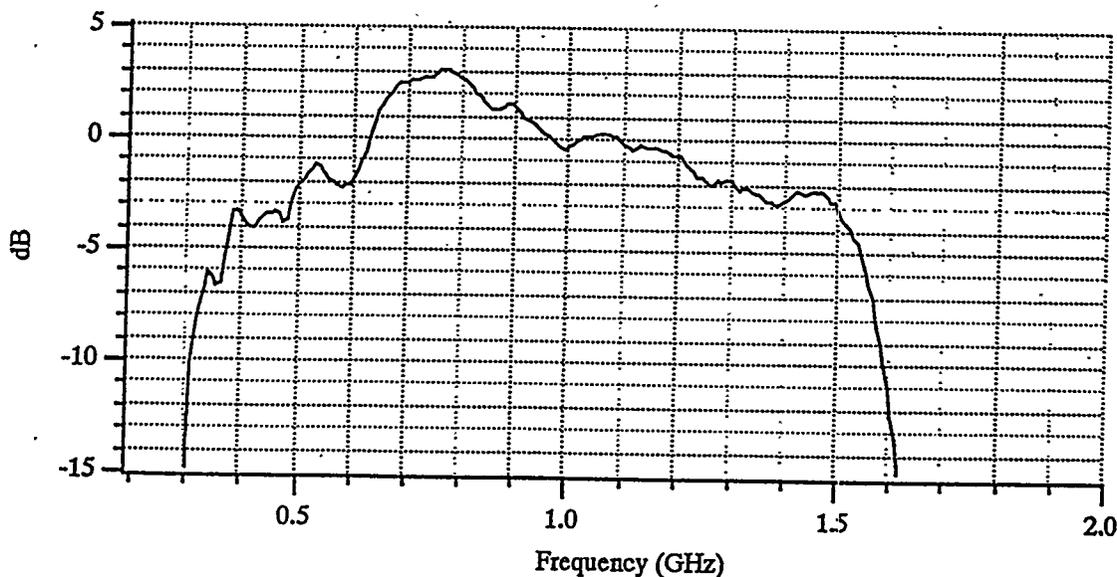


Figure 7. Normalized frequency response of antenna pair

The combined response of the transmitter, antennas, and coaxial cables is well behaved. (See Figure 8.) A frequency-domain representation of this response, shown in Figure 9, reveals that our GPR system performs well between 400 MHz and 1000 MHz. The narrower bandwidth of the monopole corner reflector acts as a filter for the horn antenna to produce this clean response. In our experience, the use of identical antennas for transmission and reception yields a response that rings.

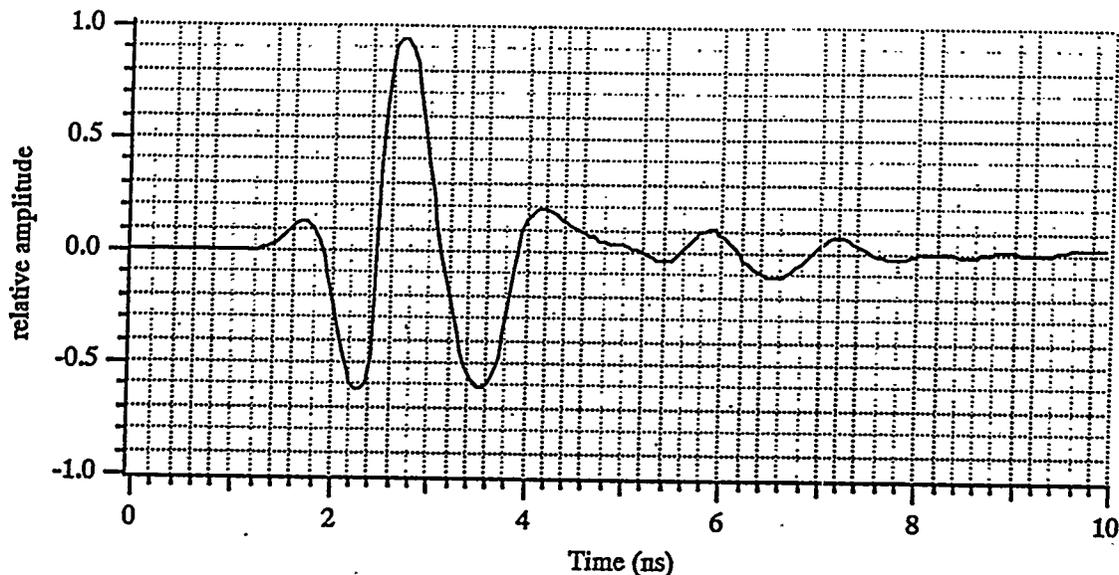


Figure 8. Time-domain response of GPR system

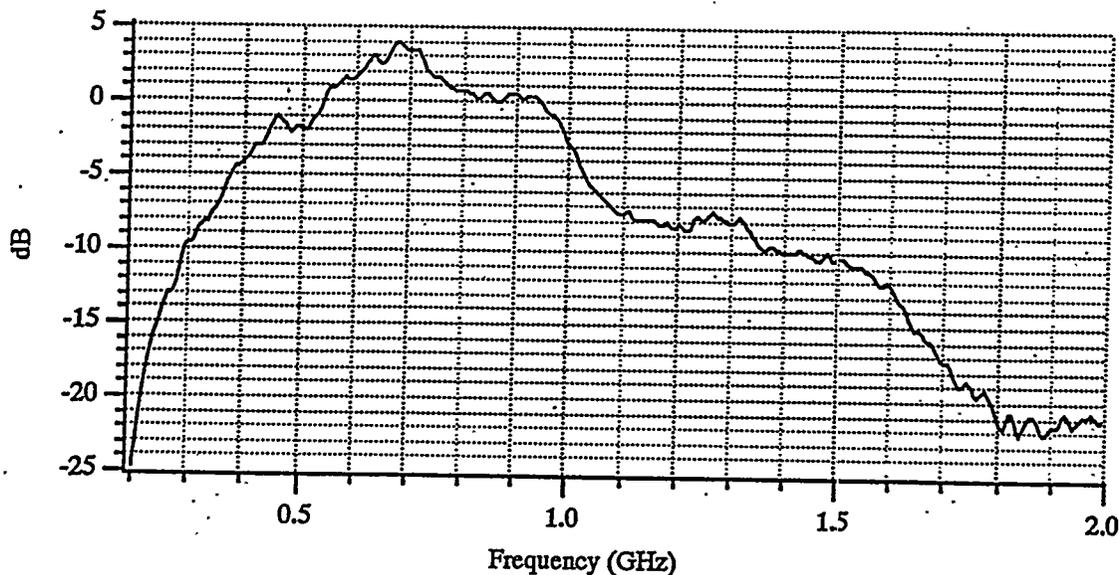


Figure 9. Normalized frequency-domain response of GPR system

Our receiver is a Hewlett-Packard 54720D 4 gigasample per second transient digitizer. We use the digitizer's internal amplifier to take advantage of its full eight-bit resolution. When triggered, the 54720D captures a 256-point waveform. Each waveform is transferred to an Apple Macintosh IIx and stored in RAM until the data run is complete. (The data acquisition and instrument control software is written in National Instruments LabVIEW, a high-level object-oriented software package.) Finally, the data is transferred to a Sun or HP workstation for SAR processing.

The antennas are mounted on a remotely-controlled pan-and-tilt mechanism, permitting adjustment of the antenna look angle. The transmitted pulse is optimally refracted into the ground when the antenna tilt angle matches the Brewster angle, as determined by the soil dielectric constant at the center of the spectrum. This is illustrated in Figure 10. Maximum penetration depth varies, depending on the conductivity of the soil.

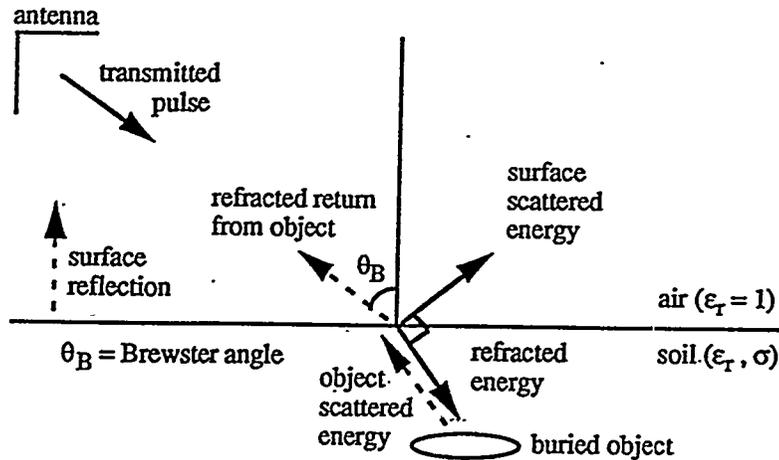


Figure 10. Illustration of rf energy propagation for buried object detection

4. THE NTS MINE DETECTION FACILITY

LLNL operates the Buried Object and Mine Detection Facility at the Nevada Test Site (NTS), which is located in the high desert, approximately 100 km northwest of Las Vegas. This facility appears in the foreground of the photograph in Figure 11. Actual mines (without detonators) and surrogate mines, both metal and plastic, have been buried in natural vegetation. The soil in this area is made up of alluvium, consisting of Paleozoic fragments and tuff. Soil conductivity is on the order of five to eight millisiemens. The exact location of buried mines has been carefully documented. Figure 12 shows how the facility is laid out. An area has been cleared of vegetation and smoothed to permit evaluation of our GPR system with a minimum of ground clutter.

Figure 13 illustrates the typical layout of a minefield plot. Items marked "fid" are 10-cm square galvanized steel plates that we use as surface markers. The fids are attached to the ground by means of short metal skewers. At each end of a minefield, three 1-cm diameter rebars are driven approximately 30 cm into the ground, leaving 60 cm above ground. Mines and surrogates are buried with between one and ten centimeters of soil overburden. The actual location of objects varies from the square grid when bushes are present. Vegetation and burrows were left intact to maintain realistic conditions.

5. EXPERIMENTAL RESULTS

We set up our GPR system at the Nevada Test Site in the configuration discussed in section 3, and positioned our antennas at an elevation of five meters. Both antennas were mounted for vertical polarization, and the antenna look angle was 34 degrees below horizontal to center the main antenna lobe over the middle row of objects. The angle from the antennas to the back row of objects was around 26 degrees. For an estimated dielectric constant of 4 at the center of the spectrum, the angle to the back row approximates the Brewster angle. Thus, Brewster's law favors the back row of objects, the antenna pattern favors the middle row, and close range favors the front row.

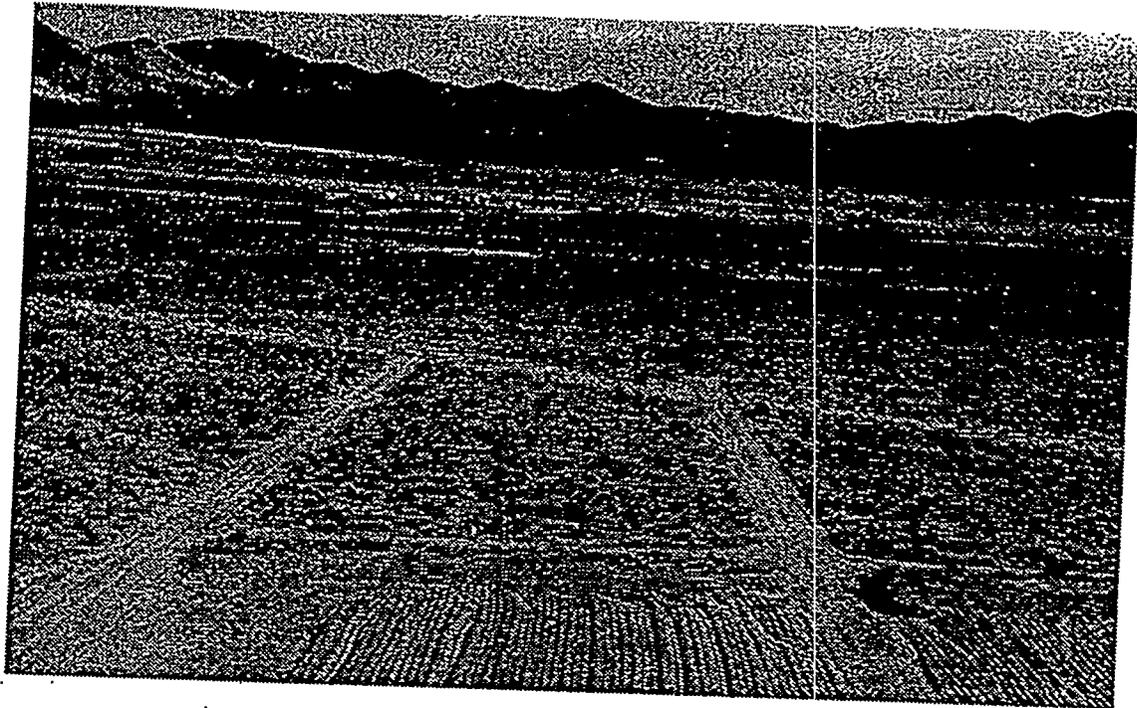


Figure 11. Buried Object and Mine Detection Facility at the Nevada Test Site

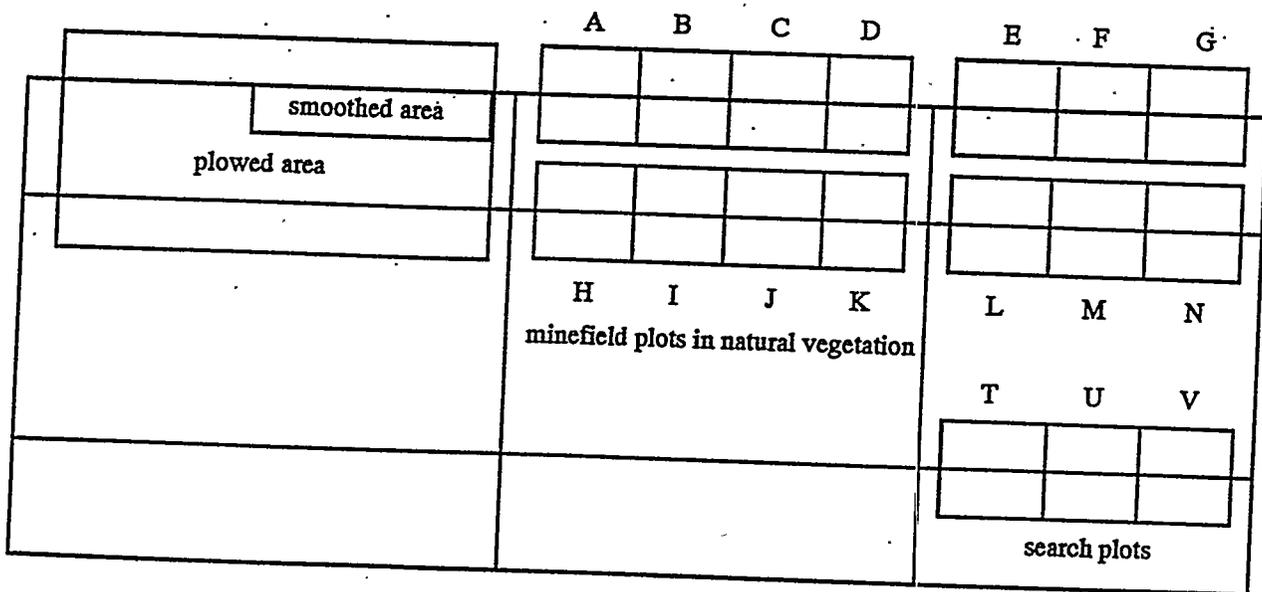


Figure 12. Layout of NTS Mine Detection Facility

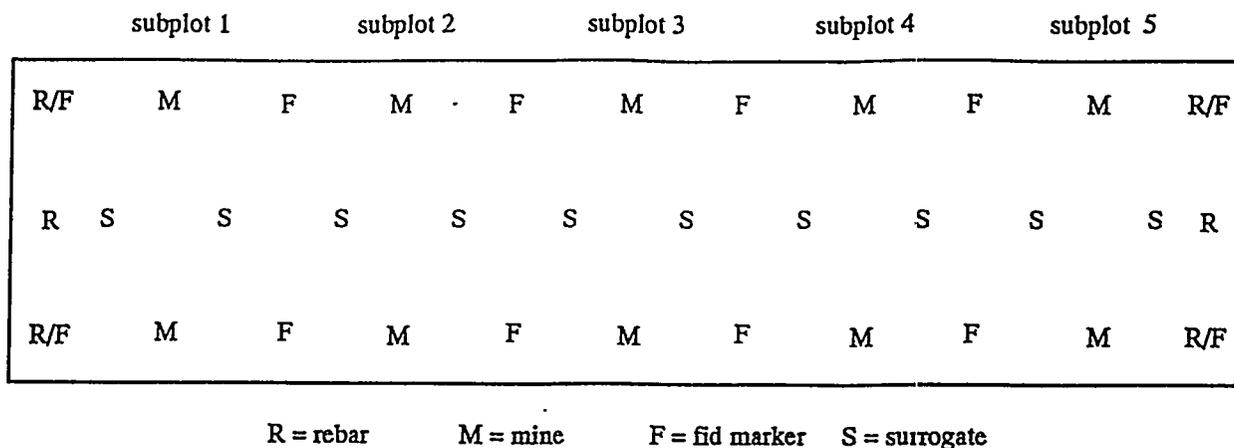


Figure 13. Typical layout of an NTS minefield plot

We collected data every five cm along the road adjacent to the minefield designated, "I-south." This minefield contains real metal mines and surrogates, all of which are approximately 30 cm in diameter. Figure 14 is a 796x256 composite image of data from subplot IS-4 with the average background subtracted. It is a side-by-side representation of the 796 time-domain data records. The three rows of objects clearly stand out above the clutter. Each object is defined by an arc. Prior to SAR processing, the image is preprocessed using average subtraction, range compensation, and pulse compression. Average subtraction takes the difference between the raw image and the mean of the raw image. Range compensation corrects for power losses due to the distance from the antennas to the objects. Pulse compression deconvolves the antenna-to-antenna system pulse response (Figure 8) from the result to reduce ringing and improve resolution.

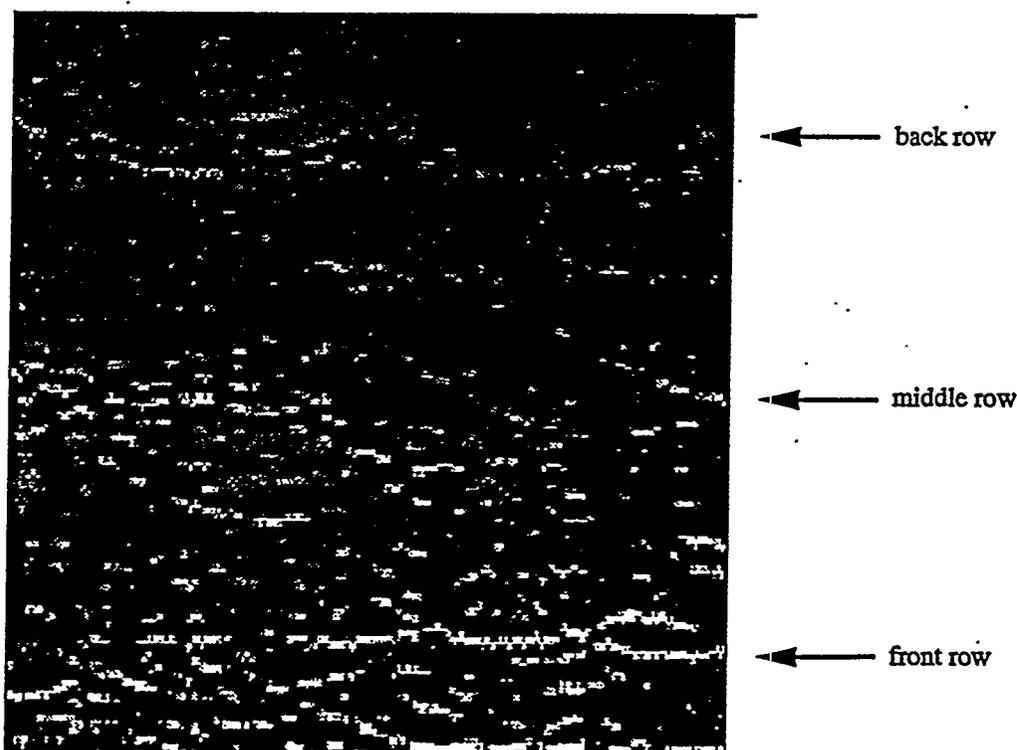


Figure 14. Composite image of minefield IS-4 subplot

Figure 15 is the reconstructed SAR image of minefield I-south, subplot-4. Object positions in this image correspond to arc positions in Figure 14. With the addition of the exact location of objects in Figure 16, we find good correlation between the

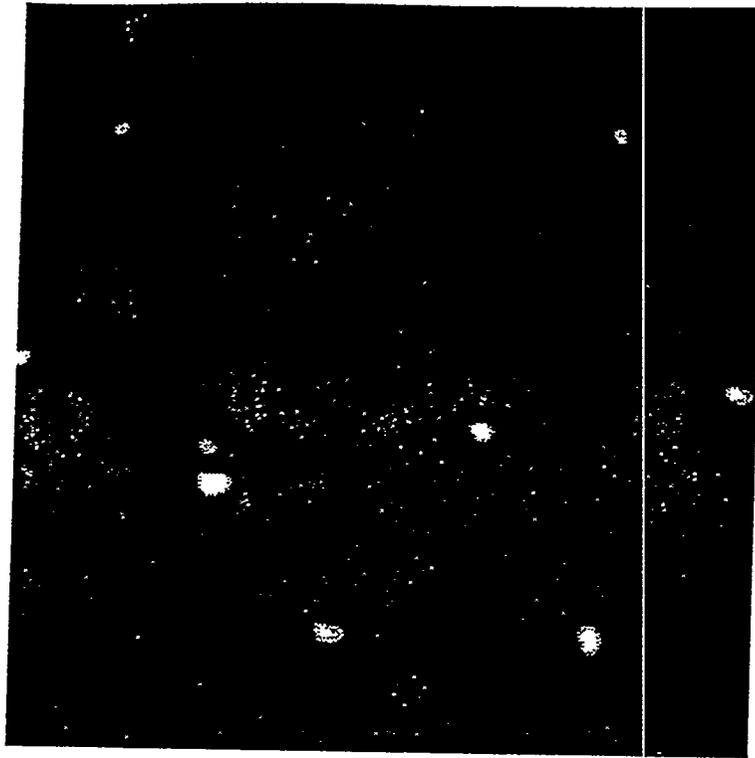


Figure 15. Reconstructed SAR image of minefield subplot IS-4

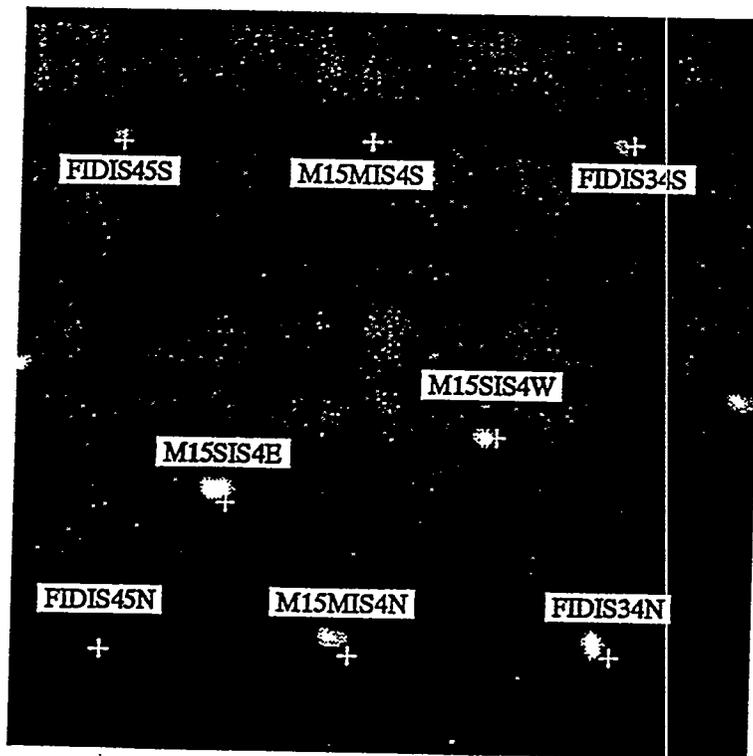


Figure 16. Reconstructed SAR image of IS-4 with ground-truth registration

ground-truth registration and the GPR result. The three-dimensional representation of this data in Figure 17 shows that most of the clutter is well below the buried objects and surface markers. By making a visual analysis of the minefield, we were able to determine that the stronger returns from clutter were due to bushes and animal burrows. Overall, the signal-to-clutter ratio in this image is between 4:1 and 6:1.

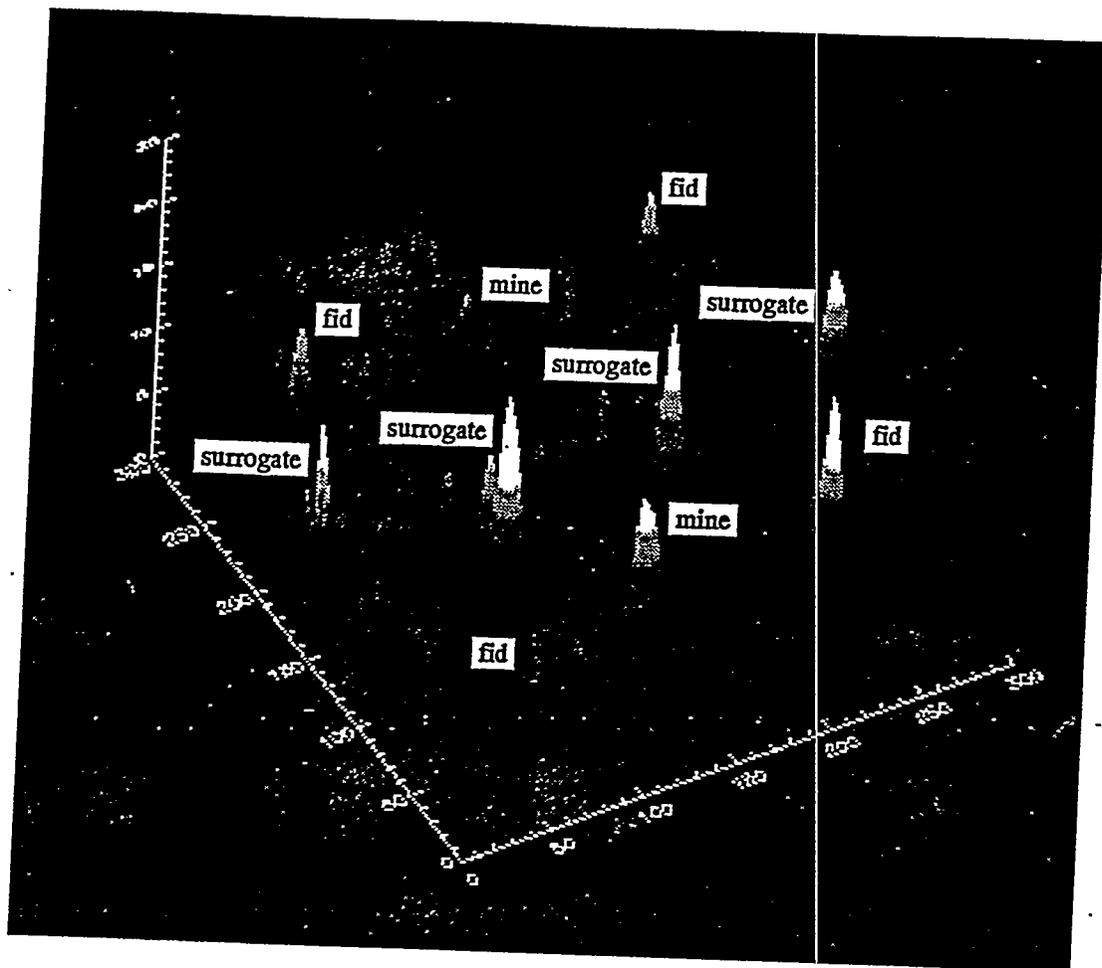


Figure 17. Three-dimensional SAR image of minefield subplot IS-4

To test the penetration capability of our GPR system, we buried five 40-cm diameter metal disks in the smoothed test plot at depths ranging from 5 cm to 40 cm. Using the same experimental setup that we used for minefield I-south, we obtained the processed result shown in Figure 18. All five of the objects are clearly visible. The three-dimensional representation of this data is shown in Figure 19. The disk with 10 cm of soil overburden produced a stronger return than the one buried 5 cm deep. It could be that the latter object is tilted back slightly, resulting in a smaller radar cross-section. The signal-to-clutter ratio for these five objects ranges between 4:1 and 6:1, depending on depth of burial.

6. SUMMARY

We have demonstrated the capability of our standoff, wideband, side-looking GPR system to locate buried metal mines with a signal-to-clutter ratio of up to 6:1 at the Mine Detection Facility at the Nevada Test Site. Reconstructed two-dimensional images of the test area compare favorably with the ground truth data, and validate the capabilities of our system. We have also demonstrated that our system can detect metal objects buried as deep as 40 cm.

The performance and functionality of our GPR system can be improved in a number of ways. We can increase the penetration depth by reducing the lower cutoff frequency of our antenna system, and by increasing the output power from our pulse transmitter. We can improve our signal-to-noise ratio by averaging multiple waveforms in each position. Modifications to our

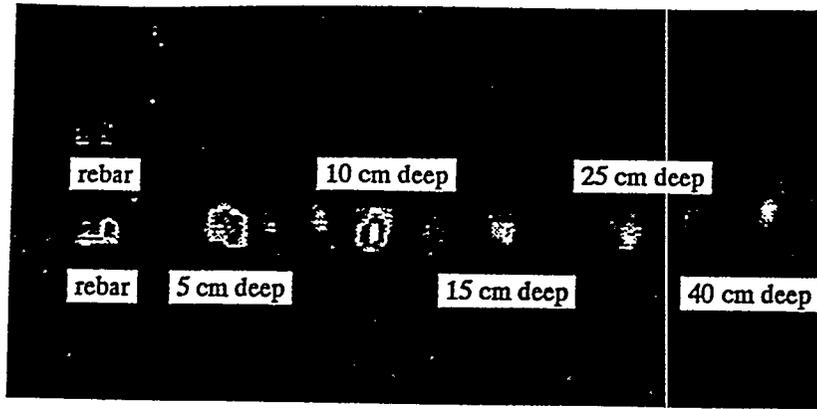


Figure 18. Reconstructed SAR image of buried metal disks

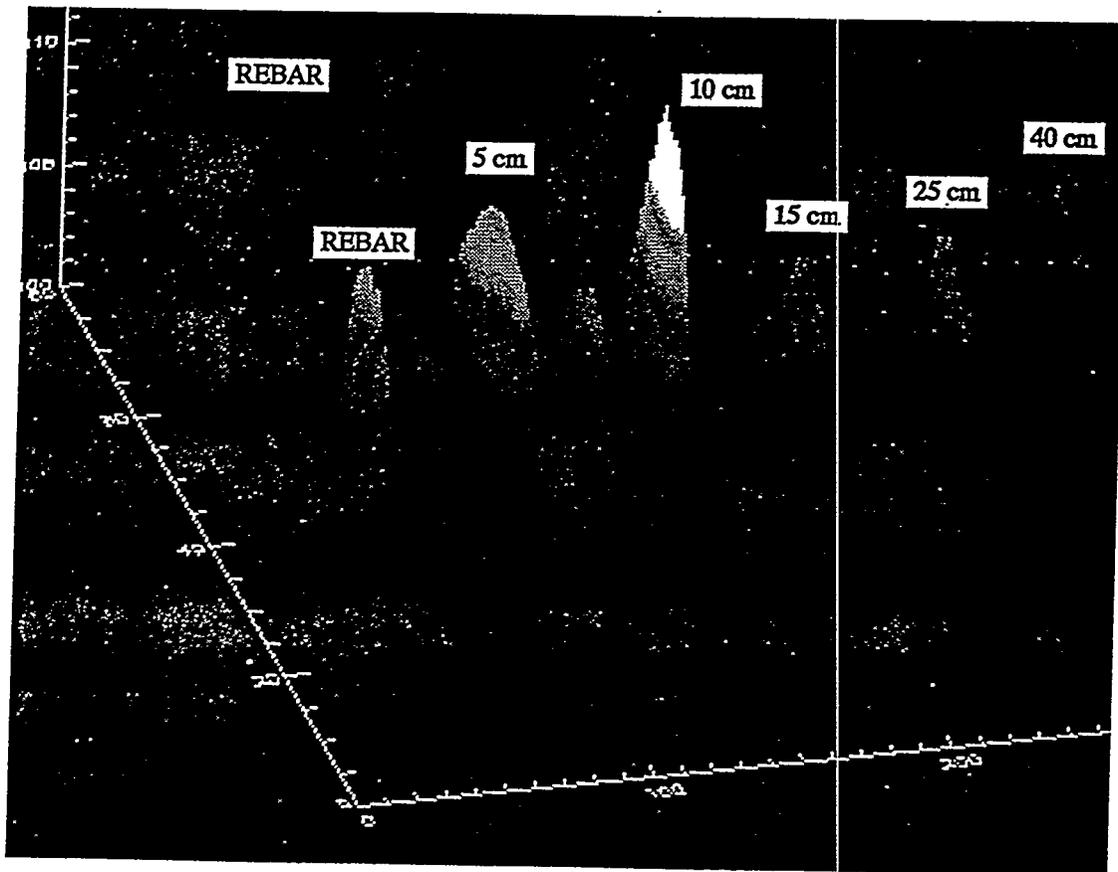


Figure 19. Three-dimensional SAR image of buried metal disks

data acquisition hardware will make it possible to place our system on a slow, low-flying airborne platform. Our intermediate goal is to eliminate the need for a boom truck by mounting our antennas on top of a motorhome. We will then be able to drive our GPR system to any test site.

7. ACKNOWLEDGMENTS

We would like to acknowledge the support of David Fields, Michael Carter, Charles Anderson, Terry Rossow, Jeff Mast, and Erik Johansson of LLNL. We also thank our field support team from EG&G, North Las Vegas.

8. REFERENCES

1. D. R. Rutt, "Use of an ultra-wideband synthetic aperture radar technique for buried mine detection," SPIE Vol. 1875, Ultrahigh Resolution Radar 1993, pp. 68-76, 20 January 1993.
2. R. S. Vickers, V. H. Gonzalez, and R. W. Ficklin, "Results from a VHF impulse synthetic aperture radar," SPIE Vol. 1631, Ultrawideband Radar 1992, pp. 219-225, 22-23 January 1992.

Work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48

Appendix B

Preliminary Results of LLNL Analysis of CSS Cobra Flight Test Data for Mine Detection

Lawrence Livermore National Laboratory

Introduction

The goal of this work is to evaluate the feasibility of using Lawrence Livermore National Laboratory's (LLNL's) automatic target recognition algorithms to detect mines in images provided by the Coastal Systems Station (CSS). The data consists of nineteen scenes, each scene consisting of six images in the wavelength bands of 400, 500, 600, 700, 800, and 900 nm. The ideal system should have a high probability of correct mine detection (P_D), and a low probability of false alarms, (P_{FA}), that is the probability of classifying an object as being a mine when it is not. Processing results show that the algorithms perform well with this data set. We expect, however, that even better performance is possible if a larger and improved data set can be obtained.

Data description and comments

The raw data images provided were 720x480 pixels in six bands from 400 to 900 nm. The active area available to us for automatic detection, however, was limited due to camera time stamps on the images. These time stamps obscured some of the mines and made it necessary for us to crop the images. The cropped image used columns 21 to 695 and rows 110 to 398 to assure that we did not get any of the text in our field of view for detection. This limited the number of targets we could use in our training set for the classifier to 133 out of a total of 209 possible objects listed in the ground truth table. The time stamps eliminated 36% of the possible targets.

The data also contained a spatial trend in the background. Background values across the image differed by up to a factor of two from one side to the other (fig1). As a result of this the signal-to-noise ratio (SNR) varied across the image and, in some cases, was quite poor. Varying SNR's cause the training sets to become less reliable and, of course, where the ratio is low, causes a decrease in contrast between mines and backgrounds. This apparently is a problem with the camera and should be calibrated out of the data. We ran a high-pass filter over the data sets to reduce this trend. This procedure is discussed below in the "Technique" section.

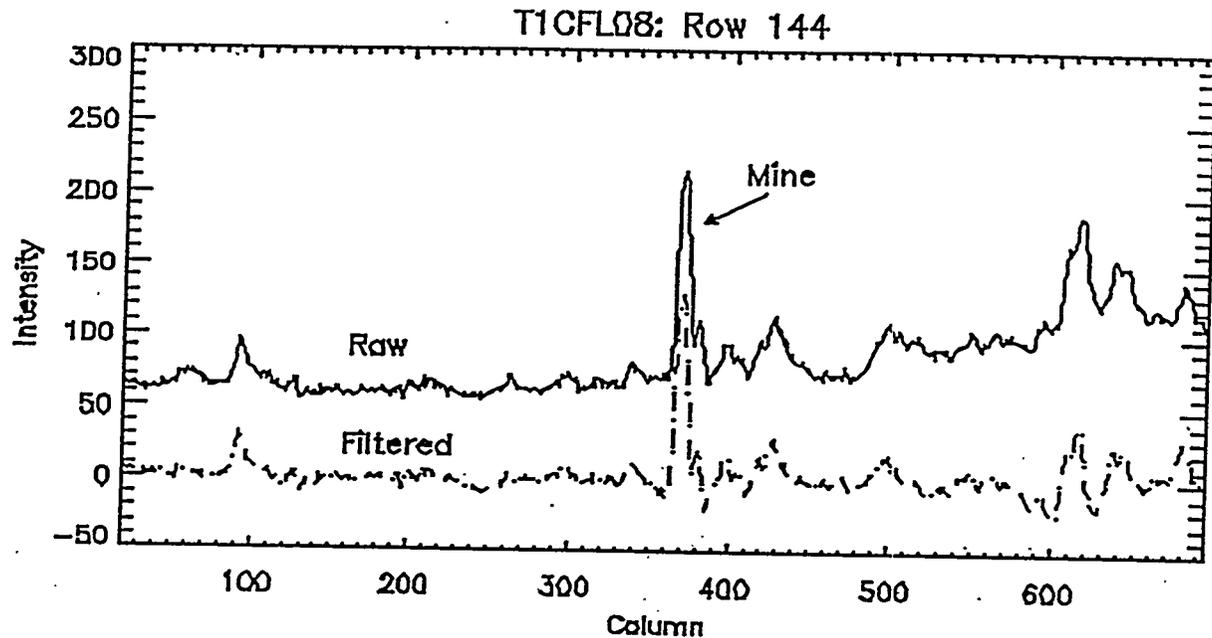


Figure 1. Row lineouts showing trend in data set

One more limitation we encountered with the data was that the resolution of the camera system was often insufficient to distinguish between circles and squares. This limits the [application of shape filters.] ability to use features based on shape for object detection. Higher resolution would allow the use of powerful techniques that distinguish shapes.

Technique

A. Preprocessing

The raw image (720x480x6 pixels) was cropped to avoid time stamps and filtering information at the top and bottom, and registration corrections on the left and right sides of the image for all bands. The range of $X=\{21:695\}$ and $Y=\{110:398\}$ was the spatial area processed. The cropped raw image contained 674x298x6 pixels.

B. Feature Extraction

The mine detection technique works by using supervised learning to train a classifier to classify each pixel in the image as belonging to one of two classes; mine or background. We create a training set for mines by cutting out 5x5x6 subimages (or "stack of tiles") containing examples of mine regions over all six bands. The spatial locations of the mine tiles are given by the ground truth provided by CSS. The background training set is obtained by cutting tiles of the same size as the mine tiles. The locations chosen for cutting background tiles, however, are randomly chosen from among the possible background areas in the scene.

Scene features are calculated from the pixels contained in the mine and background tiles. Many types of features are available, but we started with the very simple amplitude features [3,7-12, 22,23]. Amplitude features are statistical moments of the probability density function (pdf) for the pixel intensity values within the tiles. Sample moments are used to estimate the features. The features are assembled to form a feature vector, which is used as input to the classifier.

For this data, the size of the training set for mines varied from about 100 tiles (samples) for the case in which we included all types of targets, to about 50 samples for the case in which we included only metal mines as targets. The training set for the background class contained about twice as many samples as the mine class in all cases.

For this study, we used the Probabilistic Neural Network (PNN) as the classifier for this problem [6]. Because the number of available training samples is small, the classifier was trained using the "hold-one-out" technique [21].

C. Training Set

Several training sets were generated to be appropriate for the intended targets to be detected. These sets contained various types of targets: All targets, all targets without resolution panels, plastic and metal mines only, metal mines only, and plastic mines only. We searched a space of 54 features (nine features for each band) for three features selected by a Sequential Forward Selection algorithm [21,22]. Selected features differed for each type of training table. The Probabilistic Neural Net (PNN) was tuned for optimal probability of correct classification by adjusting the smoothing parameter, sigma [23]. The smoothing parameter defines the width of the Parzen window used in the PNN. The table below summarizes the type of training table, the selected features, and the optimal sigma that was obtained in each case.

training table:	band/features:	sigma:
1. All targets	400nm Max, 400nm Energy, 400nm Entropy	0.09
2. All targets, no Res panels	400nm Max, 400nm Mean, 400nm Kurtosis	0.15
3. Plastic and Metal mines	400nm Max, 900nm Mean, 500nm Max	0.80
4. Metal mines only	400nm Max, 500nm Kurtosis, 700nm Energy	0.40
5. Plastic mines only	500nm Mean, 400nm Dev, 500nm Kurtosis	0.03

D. Labeling

Once the classifier is trained, it is used to analyze an image not included in the training sets. An analysis window of the same size as the training tiles is raster-scanned over the image. At each pixel in the image, features are calculated for the pixels in the analysis window, and the PNN classifies the center pixel in the window as belong to either the class "mine" or the class "background." The resulting binary image (containing only ones representing mine pixels and zeros representing background pixels) is called the "labeled image." This labeled image provides us with an indication of the locations of probable mine pixels. The labelled image often contains "false alarm" pixels where the PNN classified the pixel as a mine pixel, when it was in fact a background pixel. When a large number of false alarms occur, their number can be greatly reduced by postprocessing the labelled image as described next.

E. Post-processing

After the trained neural net has labeled the image we apply a morphological operator to the labeled image. We use a 3x3 pixel structuring element to successively erode the labeled image, then dilate the eroded image. This operation serves to eliminate many of the small false alarms from the labeled image. The operation of erosion followed by dilation is called an 'opening.' The opening sieves out objects that are smaller than the structuring element, but avoids a general shrinking of the image [30, 34].

Since we know the physical size of our targets, we can apply size constraints to eliminate from the opened image objects that are too large or too small. First we perform a connected components analysis on the opened image [30, 34]. This operation assigns all adjacent pixels that form a region a unique number (index). We then eliminate detected object regions that are too large or too small according to prior knowledge we have about the mine size from ground truth. Given the true mine size, we discard detected regions that are smaller than 66% of the size the region is supposed to be, according to ground truth for a mine. The number "66%" was chosen based upon knowledge that the labeling process tends to erode the size of the region detected, and by some experimentation with the data set. This number can be better

tuned to the data with more work, or more sophisticated algorithms can be employed. The important thing to note is that the use of size constraints is a very useful tool for eliminating small false alarms and for using prior knowledge in the analysis.

Results and discussion

Figs. 2a - d show the resulting labeled and post-processed images of data set t1cfl08. This data set was chosen for an example because it contains a good sampling of metal and plastic mines.

We calculate the Probability of Detection (number of mine regions correctly identified / Total number of mine regions in the scene) and the Probability of False Alarm (number of background pixels identified as targets / Total number of background pixels in the scene) to assess the performance of our analysis. We consider a target correctly identified if the location of the centroid of the region is within a circular neighborhood the size of the radius of the target according to ground truth.

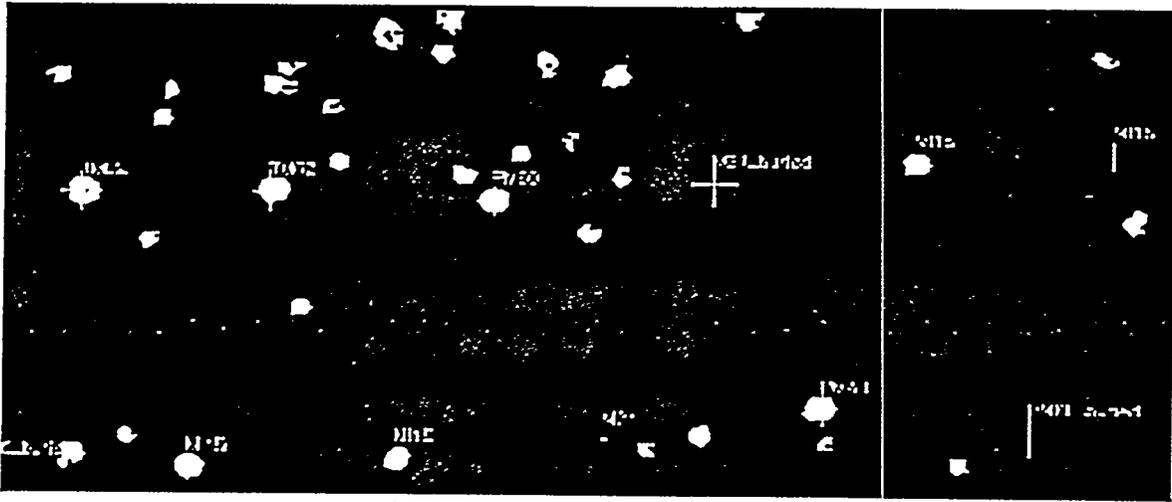


Figure 2a. T1CFL08 post processed: Trained on all targets; PD=1.00, PFA=0.0586

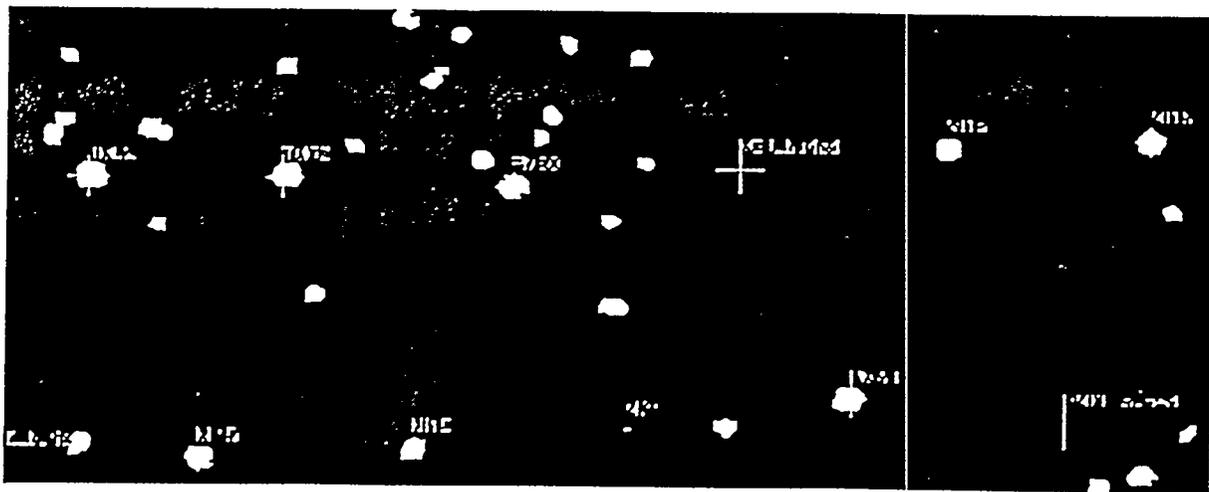


Figure 2b. T1CFL08 post processed: Trained without Resolution Panels; PD=1.00, PFA=0.036

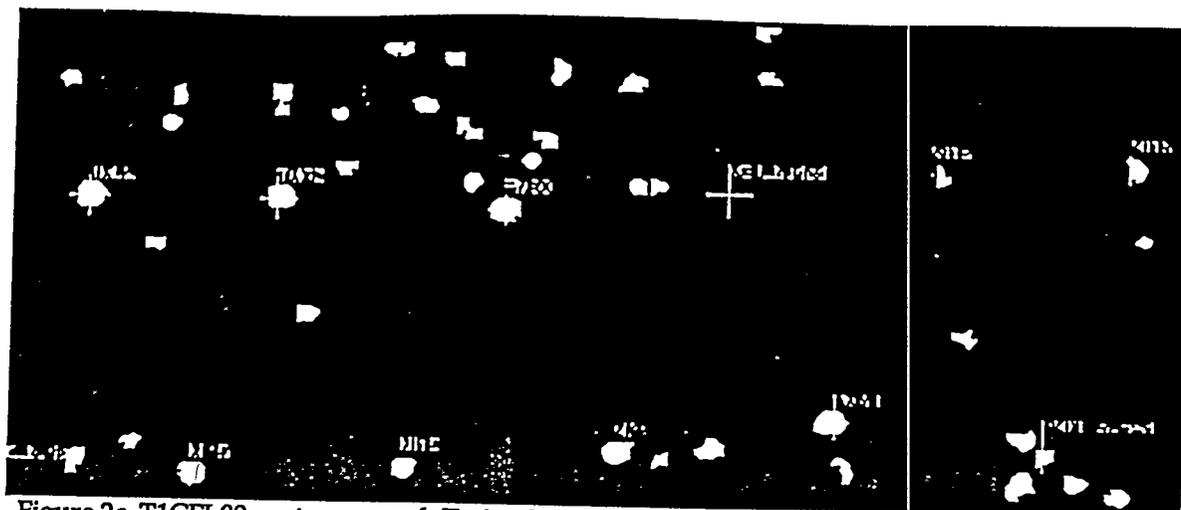


Figure 2c. T1CFL08 post processed: Trained with Metal and Plastic; PD=1.00, PFA=0.017

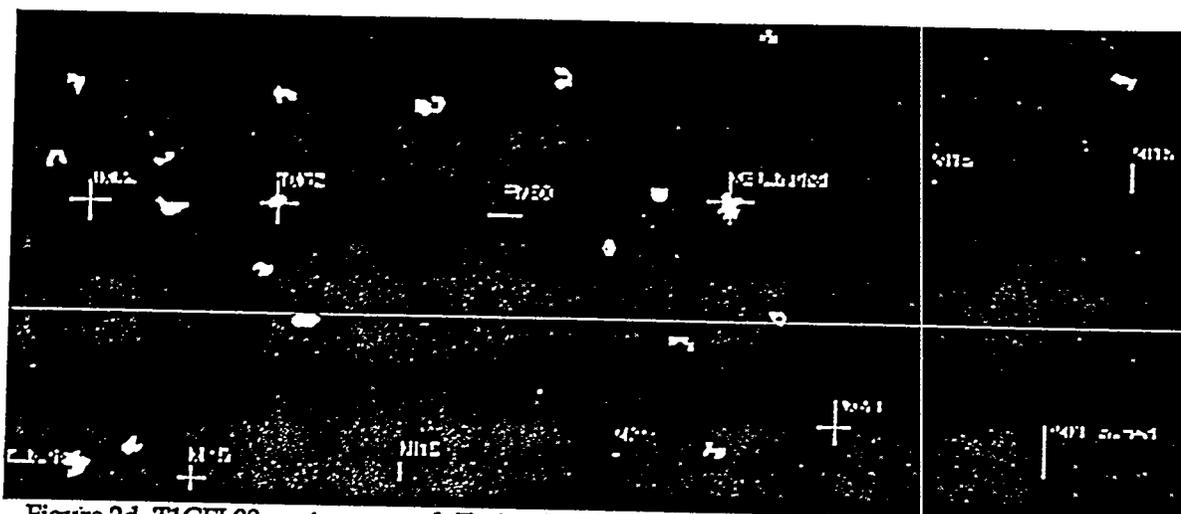


Figure 2d. T1CFL08 post processed: Trained with Plastic; PD =0.11, PFA=0.007

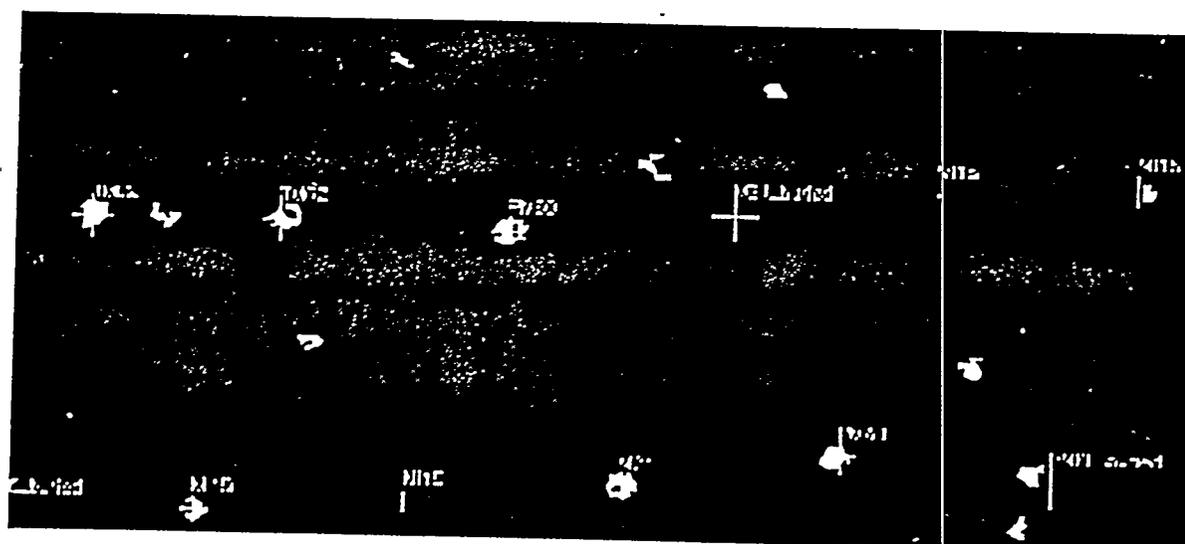


Figure 2e. T1CFL08 post processed: Trained with Metal; PD=0.77, PFA=0.0034

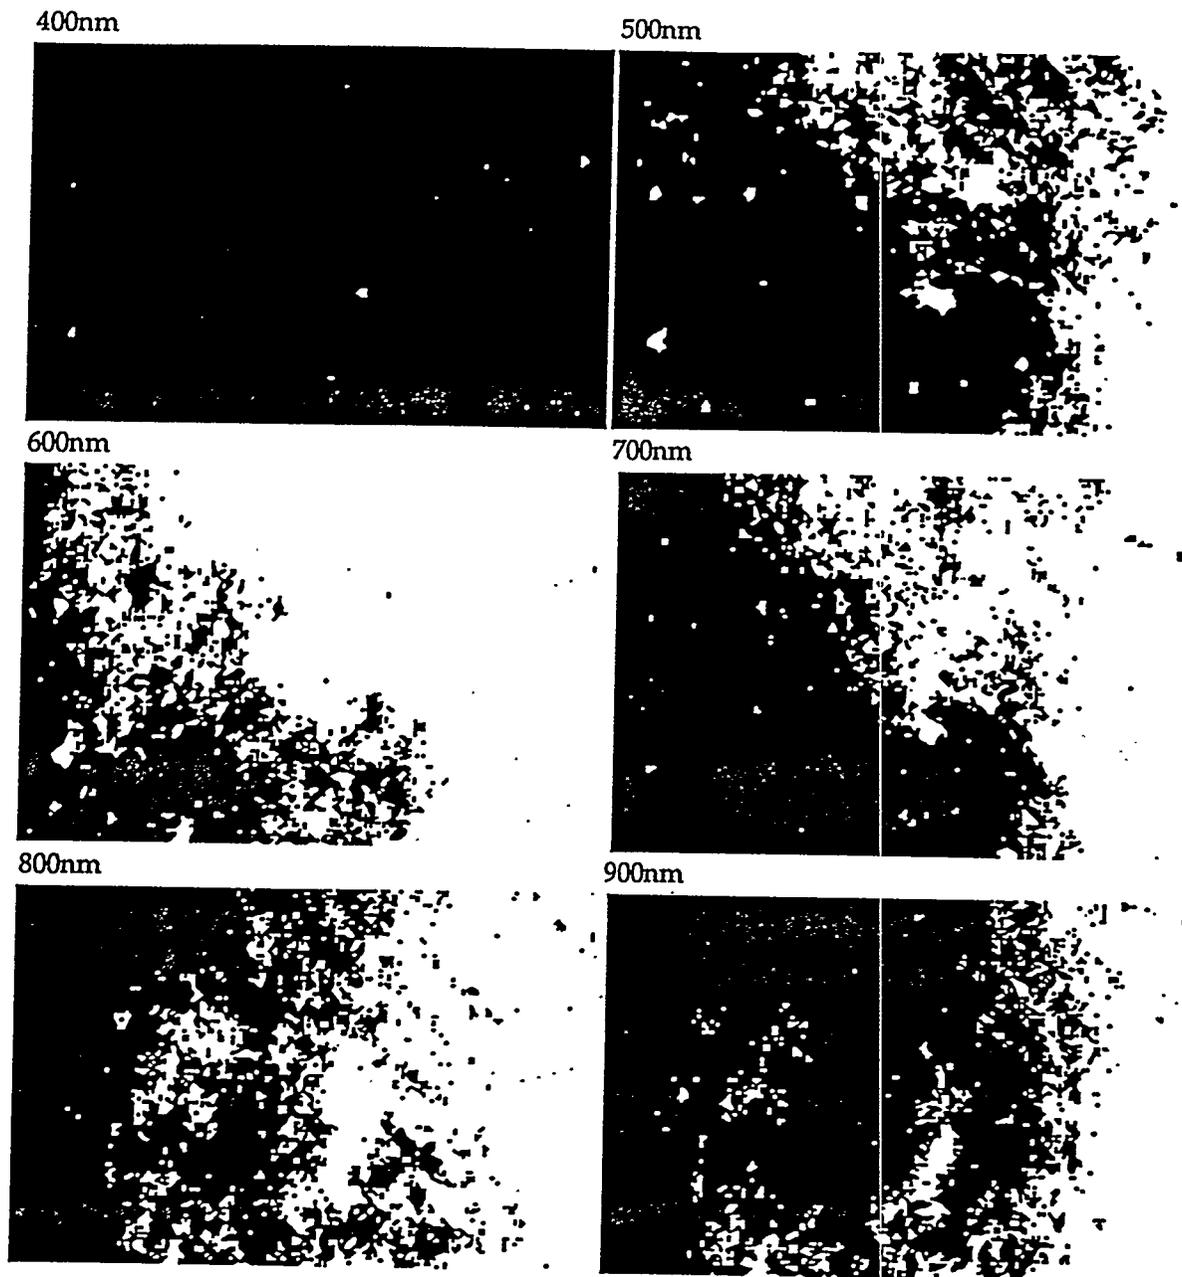


Figure 3: T1CFL08 Filtered

The cropped raw data sets were operated on with a simple highpass filter. For each pixel, the raw value was replaced by the difference of the raw value and the average of a 30 by 30 area surrounding the pixel. This differencing of a "boxcar" smoothed area greatly reduced the background trend in the raw image (fig 1).

For the case in which all objects in the ground truth were used to train the PNN (fig. 2a) we successfully detected all of the mines so that out of the objects labeled we have a probability of detection, $P_D=1$. The probability of a false alarm (P_{FA}), however, was quite high at .0586.

Training on all objects except the resolution panels (fig. 2b) gave similar results with $P_D=1$ and $P_{FA} = .036$. Training on just metal and plastic mines (fig 2c) gave $P_D=1$ and $P_{FA} = .017$. Finally training on plastic mines only gave poor results (a low P_D), with $P_D=.11$ and $P_{FA} = .0034$ (fig. 2d).

The best results occurred when we used only metal mines in the training tables (fig2e). While P_D in this example dropped to .77, P_{FA} showed a dramatic improvement at .0034. Examples of other data sets (fig 3) show that this trend in performance is consistent throughout the data sets. Data set T1CFL011 has $P_D=1$ and $P_{FA} = .0013$. T1CFL14 shows excellent results in the labeled image only (fig3b) with $P_D=1$ and $P_{FA} = .0028$. However, size constraints applied to this image produced a degradation in the results (fig 3c), with $P_D= .41$ and $P_{FA} = .00089$. Clearly, more investigation needs to be done to find the optimal size constraint that will allow the procedure to be robust with respect to all the data sets.

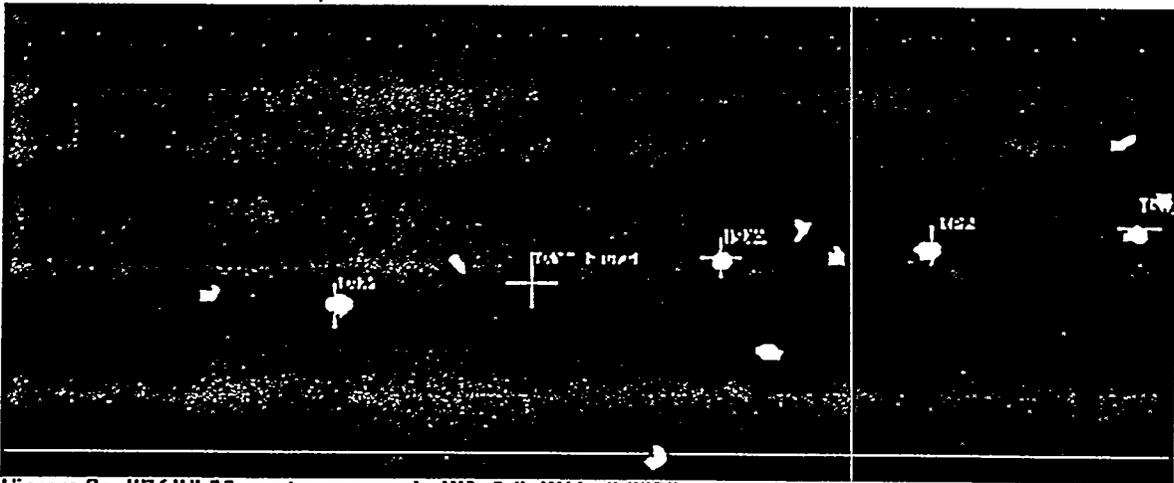


Figure 3a. T1CFL11 post processed. $P_D=1.0$, $P_{FA}=0.0030$

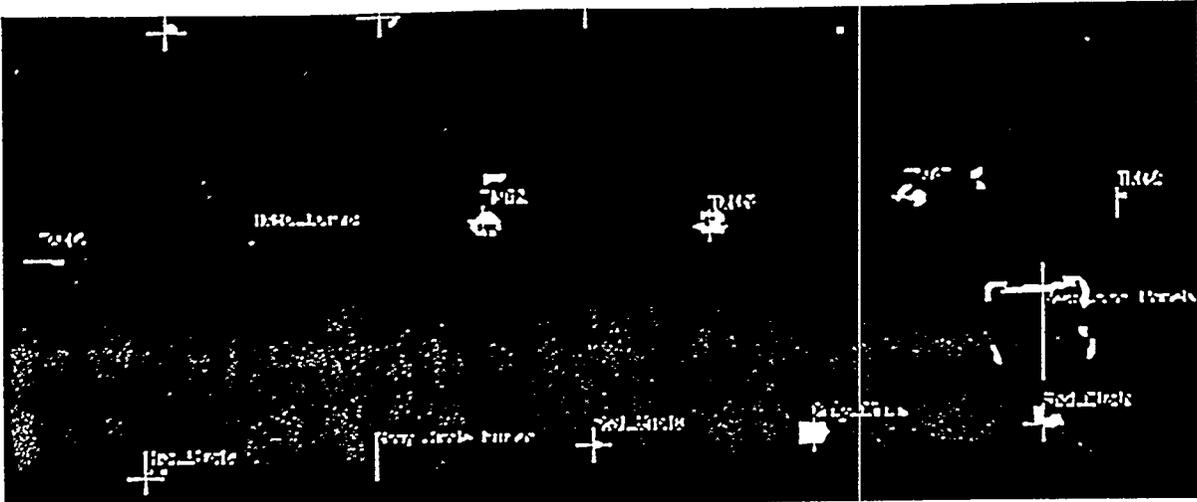


Figure 3b. TICFL14: Metal only; PD=1.0, PFA=.0028

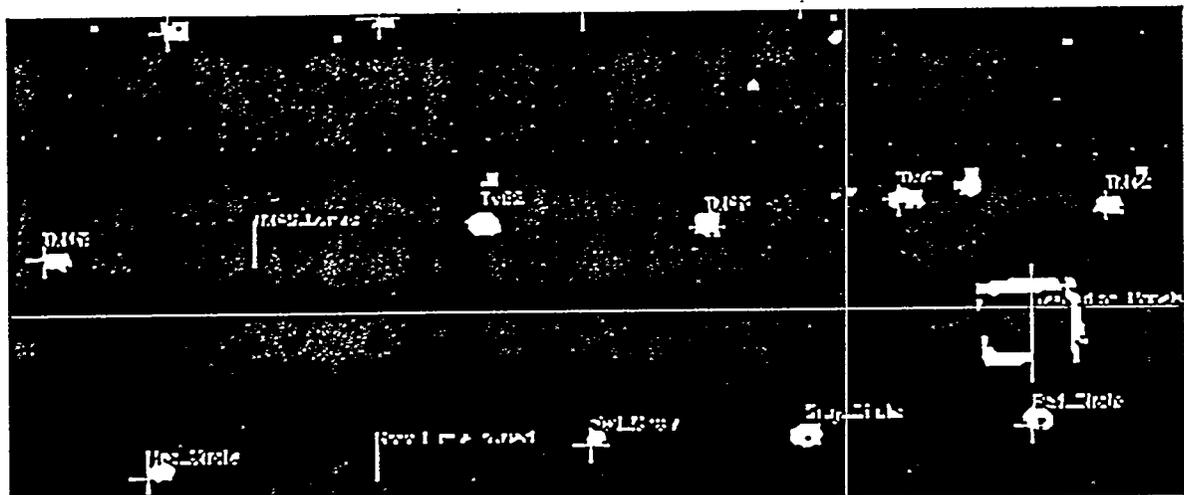


Figure 3c TICFL14: Metal only; Size constraints applied. PD=0.41, PFA=0.00089

The apparent success of the mines-only training set is most likely due to the higher contrast between metal mines and background. When training sets are used containing plastic mines the performance is diminished by the fact that in many of the bands in the data sets the plastic mines have characteristics much like the background. This is evident in fig. 2d.

Future Work

The results of the processing are very encouraging. The probabilities of detection and false alarm achieved during training and testing were sufficient for locating minefields from a qualitative visual inspection of the labeled and postprocessed images. They are very good, given

the constraints on the analysis. The current limitations, or barriers to progress, are due primarily to the following factors: (1) The data contain a relatively small number of mine training samples for the individual mine types. Of course, the principles of supervised learning rely on the assumption that the training set is sufficiently large and sufficiently representative of the test set to allow for good detection. This data set is smaller than it should be for good training. (2) The images contain a spatial trend in intensity values that is caused by problems with the camera. We removed the trend with image processing methods, but it should, ideally, be removed from the hardware. CSS is aware of this problem and corrective action is being taken. (3) More sophisticated methods that make use of shape information could be used if the camera resolution were greater. If the resolution is fixed for all possible scenarios, this issue may be moot. (4) Only a very simple analysis was attempted, due to funding and time constraints. Further work involving use of more powerful algorithms promises to provide improved results.

Performance improvements are expected if we address these current limitations as follows: (1) Increasing the number of training samples by conducting new experiments. (2) Solving the problems associated with the camera, (spatial trend and resolution limitations), (3) Conducting a more thorough study of the features and of which ones add the most value to the analysis. To date, we have used only amplitude features and a suboptimal but fast and efficient feature selection algorithm. A better study can be done using our optimal algorithm. We can also apply more physical knowledge and human judgment to the feature selection process to provide a more thorough analysis. (4) Tuning the tile size for optimum performance, (5) Using shape-based features on higher-resolution images. (6)

Conclusion

The algorithms developed at LLNL for buried land mine detection have been applied to the CSS flight test data containing surface land mines. The algorithm performance is good and quite encouraging. Probability of detection is very high (often close or equal to one). The probability of false alarm is low enough to allow qualitative detection of a minefield from visual inspection of the labeled and postprocessed images. The main issue of concern is that some labeled images are corrupted by small false alarm regions. Many of these are eliminated by applying size constraints, and we expect that many more could be eliminated by shape constraints. The current barriers to progress are small sample size, relatively low camera resolution, spatial trends in camera intensity, and fiscal constraints. We expect that the performance can be improved significantly by addressing these current limitations in future work.

Acknowledgments

The authors gratefully acknowledge our sponsors; David Vaughn of the Marine Corp., Quantico, VA, Ned Witherspoon and Ab Dubey of CSS, who also provided the image data, and Mark Eckart (LLNL), who served as LLNL co-sponsor. We thank Michael Carter and Charles Anderson of LLNL, who provided valuable input for the analysis.

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. G. A. Clark, Jose E. Hernandez, Sailes K. Sengupta, Robert J. Sherwood, Paul C. Schaich, Michael R. Buhl, Ronald J. Kane, Marvin J. Barth and Nancy K. DelGrande, "Computer Vision and Sensor Fusion for Detecting Buried Objects," SPIE's 1993 International Symposium and Exhibition on Optical Engineering and Photonics in Aerospace and Remote Sensing, Conference on Underground and Obscured Imaging and Detection, Orlando, FLA, April 12-16, 1993.
8. G. 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, "Sensor Feature Fusion for Detecting Buried Objects," SPIE's 1993 International Symposium and Exhibition on Optical Engineering and Photonics in Aerospace and Remote Sensing, Conference on Underground and Obscured Imaging and Detection, Orlando, FLA, April 12-16, 1993.
9. Michael R. Buhl, Jose E. Hernandez, Gregory A. Clark, and Sailes K. Sengupta, "Dual-Band Infrared Buried Mine Detection Using a Statistical Pattern Recognition Approach," Lawrence Livermore National Laboratory report UCRL-ID-114838, August, 1993.
10. G. A. Clark, S. K. Sengupta, M. R. Buhl, R. J. Sherwood, P. C. Schaich, N. Bull, R. J. Kane, and M. J. Barth, "Detecting Buried Objects by Fusing

Dual-Band Infrared Images, "*Invited paper*," Asilomar Conference on Signals, Systems and Computers, Pacific Grove, CA, Nov. 1-3, 1993.

11. N. K. DelGrande, P. F. Durbin, D. E. Perkins, G. A. Clark, R. J. Sherwood, J. E. Hernandez, and A. B. Shapiro, "Dual-Band Infrared Capabilities for Imaging Buried Object Sites," SPIE's 1993 International Symposium and Exhibition on Optical Engineering and Photonics in Aerospace and Remote Sensing, Conference on Underground and Obscured Imaging and Detection, Orlando, Florida, April 12-16, 1993.
12. M. R. Buhl, J. E. Hernandez, S. K. Sengupta, G. A. Clark, "Dual-Band Infrared Buried Mine Detection Using a Statistical Pattern Recognition Approach," Lawrence Livermore National Laboratory report UCRL-114838, August, 1993.
13. G. A. Clark, "Report on the Workshop on Automatic Target Recognition (ATR) for Airborne Standoff Mine Detection (ASMD)," Sponsored by DARPA and the U. S. Marines at the System Planning Corporation, Arlington, VA (Workshop: December 9-10, 1992), March 5, 1993.
14. G. A. Clark (Workshop Chairman), J. E. Hernandez, S. K. Sengupta, R. J. Sherwood, M. R. Buhl, P. C. Schaich, R. J. Kane, M. J. Barth, and N. K. DelGrande, "Computer Vision and Sensor Fusion for Detecting Buried Mines," Workshop on Automatic Target Recognition (ATR) for Airborne Standoff Mine Detection (ASMD), Sponsored by DARPA and the U. S. Marines at the System Planning Corporation, Arlington, VA, December 9-10, 1992.
15. G. A. Clark, J. E. Hernandez, S. K. Sengupta, R. J. Sherwood, P. C. Schaich, M.R. Buhl, R. J. Kane, M.J. Barth, and N. K. DelGrande, "Computer Vision and Sensor Fusion for Detecting Buried Objects," Asilomar Conference on Signals, Systems and Computers, Pacific Grove, CA, October 26-28, 1992.
16. 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.
17. 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.

18. "Airborne Detection of Buried Minefields," Energy and Technology Review, Lawrence Livermore National Laboratory, December, 1991.
19. 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.
20. Hu, M.K., 1962. "Visual Pattern Recognition by Moment Invariants", IRE Trans. Info. Theory, vol. IT-8, pp. 179-187.
21. Devijver, P. A. and J. Kittler, 1982. Pattern Recognition: A Statistical Approach. Prentice Hall, Englewood Cliffs, 448 pp.
22. 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.
23. Specht, D., 1990b. "PNN and Polynomial Adaline as Complementary Techniques for Classification", IEEE Trans. Neural Networks, vol. 1, pp 111-121.
24. Cacoullos, T., 1966. "Estimation of a Multivariate Density", Annals of the Institute of Statistical Mathematics (Tokyo), vol. 18(2), pp. 179-198.
25. Murthy, V.K. 1965. "Estimation of Probability Density", Annals of Mathematical Statistics, vol. 36, pp. 1027-1031.
26. 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.
27. Parzen, E., 1962. "On Estimation of a Probability Density Function and Mode", Annals of Mathematical Statistics, vol. 33, pp. 1065-1076.
28. 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.
29. Hand, D. J., 1981. Discrimination and Classification, New York, L. Wiley and Sons, pp. 218.
30. A. K. Jain, Fundamentals of Digital Image Processing, Prentice Hall, 1989.
31. Gonzales, Rafael C. and Wintz, Paul., Digital Image Processing, second

edition. Addison-Wesley Publishing Company, 1987.

32. Hall, Ernest L., Computer Image Processing and Recognition. Academic Press, 1979.
33. Rosenfeld, Azriel and Kak, Avinash C, Digital Picture Processing, Academic Press, 1976.
34. R. M. Haralick and L. G. Shapiro, Computer and Robot Vision, Vol. 1, Addison Wesley, 1992.