Pacific Northwest National Laboratory

Operated by Battelle for the U.S. Department of Energy

Advanced Image Collection, Information Extraction, and Change Detection in Support of NN-20 Broad Area Search and Analysis

G. M. Petrie
E. M. Perry
R. R. Kirkham
D. E. Slator
P. G. Heasler
H. P. Foote

September 1997

Prepared for the U.S. Department of Energy
under Contract DE-AC06-76RLO 1830

PNNL-11729

DISCLAIMER

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC06-76RLO 1830

Printed in the United States of America

Available to DOE and DOE contractors from the Office of Scientific and Technical Information, P.O. Box 62, Oak Ridge, TN 37831; prices available from (615) 576-8401.

Available to the public from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Rd., Springfield, VA 22161


This document was printed on recycled paper.
(9/97)

PNNL-11729

Advanced Image Collection, Information Extraction, and Change Detection in Support of NN-20 Broad Area Search and Analysis

G. M. Petrie
E. M. Perry
R. R. Kirkham
D. E. Slator
P. G. Heasler
H. P. Foote

September 1997

Prepared for
the U.S. Department of Energy
under Contract DE-AC06-76RLO 1830

Pacific Northwest National Laboratory
Richland, Washington 99352

Executive Summary

This report describes the work performed at the Pacific Northwest National Laboratory (PNNL) for the U.S. Department of Energy's Office of Nonproliferation and National Security, Office of Research and Development (NN-20). The work supports the NN-20 Broad Area Search and Analysis, a program initiated by NN-20 to improve the detection and classification of undeclared weapons facilities.

Ongoing PNNL research activities are described in three main components: image collection, information processing, and change analysis. The Multispectral Airborne Imaging System, which was developed to collect georeferenced imagery in the visible through infrared regions of the spectrum, and flown on a light aircraft platform, will supply current land use conditions. The image information extraction software (dynamic clustering and end-member extraction) uses imagery, like the multispectral data collected by the PNNL multispectral system, to efficiently generate landcover information. The advanced change detection uses a priori (benchmark) information, current landcover conditions, and user-supplied rules to rank suspect areas by probable risk of undeclared facilities or proliferation activities. These components, both separately and combined, provide important tools for improving the detection of undeclared facilities.

Contents

Figures

1.0 Introduction

The research described in this report was performed at the Pacific Northwest National Laboratory (PNNL) through direction and funding by the U.S. Department of Energy's Office of Nonproliferation and National Security, Office of Research and Development (NN-20). This office has the responsibility of guiding the development of technologies for detecting and identifying facilities involved in the development, assembly, and/or testing of nuclear, chemical, or biological weapons of mass destruction. The DOE created the program called Broad Area Search and Analysis (BASA) to address the difficulties associated with the detection of such weapons facilities. One of the important technology areas BASA has focused on is the use of remote imaging for identification and monitoring of proliferation facilities; this report describes some key technologies that, independently and as a whole, contribute to the ability to detect and assess undeclared facilities.

The technologies described in this report fall into three main areas. First, a multispectral airborne imaging system is described; this system generates images in the visible, near-infrared and infrared (IR) spectral regions. The system utilizes a tunable filter for the visible wavelengths and global positioning satellite (GPS) for image georeferencing. This system's output is high-spatial and spectral-resolution imagery that can provide current landcover information.

The second component described is image information extraction software for efficiently generating landcover information. This software includes improved classification (dynamic clustering) and end member extraction that can enhance image spectroscopy approaches such as linear unmixing. In this section, a landcover classification of a nuclear facility is presented to demonstrate the challenges using current imaging.

The third component of this report is a description of PNNL's advanced change detection algorithm that uses a priori (benchmark) information, current landcover conditions, and user-supplied rules to rank areas by probable risk for the identification of undeclared facilities or proliferation activities.

The report is organized as three main sections, each outlining these three technologies. This document is intended for a broad audience with an interest in technologies to assist nonproliferation detection efforts. The report is designed in its entirety as documentation for FY97 research activities at PNNL in the area of nonproliferation monitoring, or to be read as stand-alone sections outlining specific imaging technologies.

2.0 Multispectral Airborne Imaging System

PNNL's goal in the development of an airborne imaging system has been to develop a low-cost, multispectral imaging system that can 1) operate throughout significant portions of the visible and infrared (IR) spectral range, 2) operate reliably from low-cost light aircraft, 3) perform image acquisition of the required accuracy and dynamic range to detect features associated with undeclared facilities, and 4) be sufficiently flexible to allow real-time variation of operating procedures. To the extent possible, commercial off-the-shelf components have been used in the design and construction of the system. A fully integrated, self-contained IR camera and associated electronics provides quantitative thermal imagery in the 3-5 micron spectral range. A unique tunable filter provides scene measurement at specific spectral bands throughout the visible portion of the spectrum. The specific bands selected may be tailored in real time to optimize the system for feature identification and change detection within the scene being investigated. The following is a discussion of design requirements and the resulting hardware and software configuration.

2.1 Hardware Description

The current imaging system consists of a filtered charge coupled device (CCD) camera for detecting signals with wavelengths in the visible portion of the spectrum, a focal plane array sensitive to emissions in the mid-IR region, and a commercial digital camcorder for directional sighting and archival purposes. This array is shown in Figure 2.1. Each of these components will be discussed, together with data gathered during the initial flight testing program.

2.1.1 Visible Imaging System

The high-resolution and wide dynamic range provided by recent progress in CCD technology has made this the detector of choice for this system. CCD's are silicon solid-state devices that basically convert incident photons into a proportional quantity of electrical charge and store this charge within a MOS capacitor. This charge is then sequentially transferred to an on-chip amplifier and associated analog-to-digital converter. In many cases, the large range of illumination is encountered within the same image, driving the need for wide dynamic range. Since the dynamic range can be thought of as the ratio of pixel saturation to noise, it has been necessary to use a cooled, slow-scanned scientific-grade sensor to reduce system noise. The excellent linearity characteristics of current CCD's throughout the entire dynamic range provides a linear signal transfer relationship to within a fraction of 1%.

To provide for sufficiently high spatial resolution, a 1317 x 1035 pixel imaging array was chosen to incorporate into a vacuum housing. For the imaging system to achieve high dynamic range, it is critical that the sensor be cooled to a -25°C operating temperature within the vacuum chamber by using a multistage thermoelectric cooler. The fine pitch of the individual pixels, 6.8 microns, allows for high-speed readout; the 12-digit digitization at 5 megapixel/second readout rate yields a full-frame readout in 0.3 seconds.

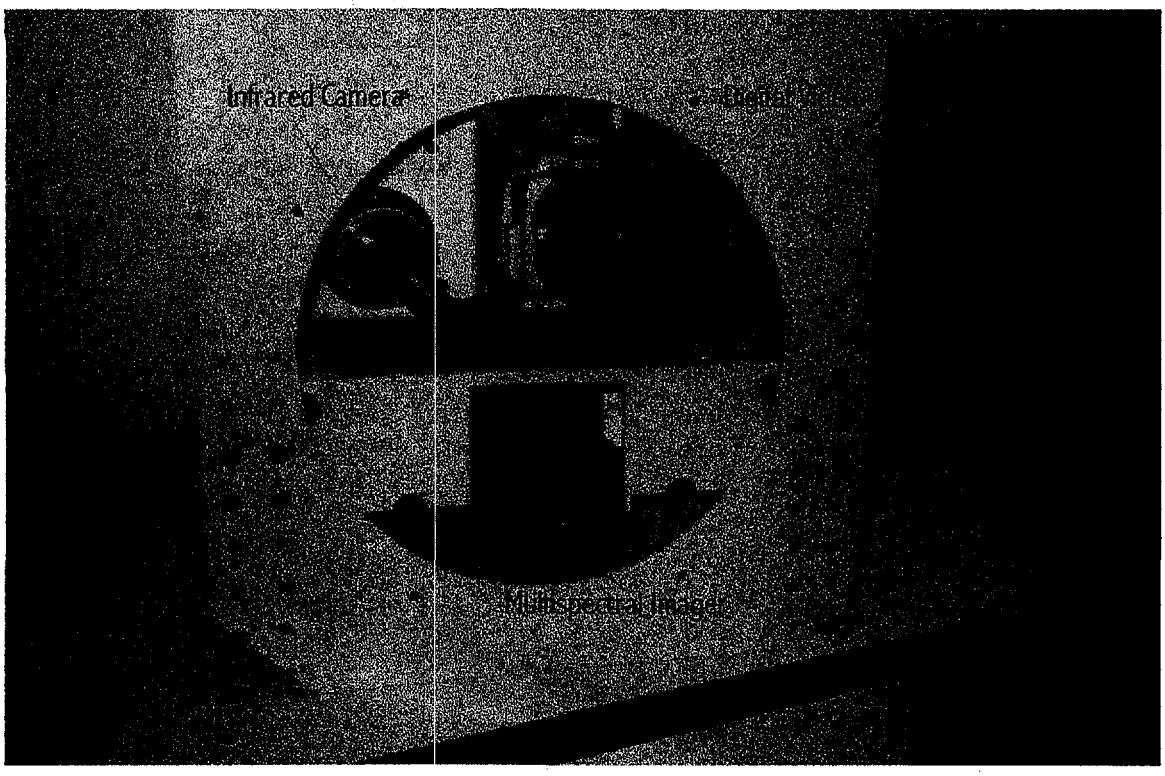

Figure 2.1. The Multispectral Airborne Imaging System with Front Windowing Plate Removed. Images coming from the three separate cameras are available for image processing, archival, and viewing by the operator at the accompanying control computers.

The CCD is a full-frame device, with exposure times determined by a conventional electromechanical shutter that is controlled by the master computer system command sequence. Preceding the shutter and CCD is an interchangeable camera lens that is selected by the user to satisfy the field of view requirements of the mission.

Spectral discrimination is required to differentiate the wavelengths of interest from a background rich in spectral signatures. The final spectral range available to the multispectral, visible imaging system is determined by the optical properties of silicon, which is opaque to photons with wavelengths of less than about 400 nanometers (nm) and transparent to wavelengths longer than 1100 nm and the available spectral range of the filter element, which is 400-720 nm. The spectral discrimination requirement is met by using an image-quality tunable birefringent filter that consists of several stages of linear parallel polarizers surrounding liquid crystal retarder elements, referred to as a liquid crystal tunable filter (LCTF). The LCTF includes an electronics module that provides direct computer control of wavelength selection. This filter is positioned within the optical path directly preceding the camera lens. The LCTF rapidly becomes transparent to light beyond 750 nm, leading to the requirement that the final component of the optical path be a "hot mirror," specifically rejecting light with wavelengths greater than 750 nm that would otherwise be registered by the CCD sensor as being of the wavelength selected by the LCTF. The passband is randomly

selectable throughout the 400-720 nm range, requiring a few tens of milliseconds for wavelength selection and stabilization. Although the spectral position of the passband is fully controlled by the imaging system user, the bandwidth is dictated by the mechanics of the LCTF; that is, it is constant in frequency space ($\Delta\nu/\nu$ is constant), and for the present system, is nominally 10 nm at 550 nm.

2.1.2 Infrared Imaging System

The second major component of the multispectral imaging system consists of a fully integrated, self-contained IR camera and associated electronics to acquire, digitize, and store images of mid-IR scenes. The core of the camera is a cryogenically cooled 256 x 256 pixel indium antimonide (InSb) image plane sensor used in a staring mode. InSb offers the widest available dynamic range and sensitivity among current IR sensing arrays. The camera's sensitivity extends throughout the 3-5 micron spectral range. This camera is bore-sighted with the filtered CCD system to provide simultaneous, multispectral imagery of the target. The 38-micron pixel size, combined with the current 50-mm germanium-silicon lens, yields a field of view of 11 degrees, comparable to that of the filtered CCD imaging system. During operation, the InSb sensor temperature is maintained at 77°K by a closed-cycle sterling cooler. An internal temperature-controlled source calibrates the camera and normalizes the camera response for non-uniform pixel gain, DC offset variability, and non-responsive pixel sites. Following a 12-bit analog-to-digital conversion, the signal is downloaded to a high-speed digital frame grabber and image processing board within the master computer for further processing and storage. Imaging speeds are available from video rates (60 frames/sec) to individual, triggered acquisitions timed to coincide with data gathered by the filtered CCD imaging system. NTSC composite and S-video output is available to provide continual video output to the operator and may be recorded for further analysis. The interactions of the various components of the imaging system are illustrated in Figure 2.2, which also shows the data and command flows.

2.1.3 Navigational System

Post-flight analysis of data collected during operations is aided by a positional and attitude logging system incorporated onto the imaging system framework. The navigation system consists of a Global Positioning Satellite (GPS) receiver, an electronic compass, and a 6-axis dynamic measurement unit (DMU). The GPS receiver can be connected to a radio receiver to provide real-time differential GPS data (<5-meter resolution). An electronic compass utilizing magnetometers provides heading information. The compass also includes two tilt sensors that provide roll and pitch information. The DMU provides X,Y,Z acceleration channels together with roll, pitch, and yaw angular rate gyro outputs. A Campbell Scientific Inc. data logger collects data from these three sensors at 1-second intervals. The data are stored in memory for later use in Geographic Information System (GIS) software. The flight path for a June 24, 1997 test flight is shown in Figure 2.3, with flight duration increasing from light to dark.

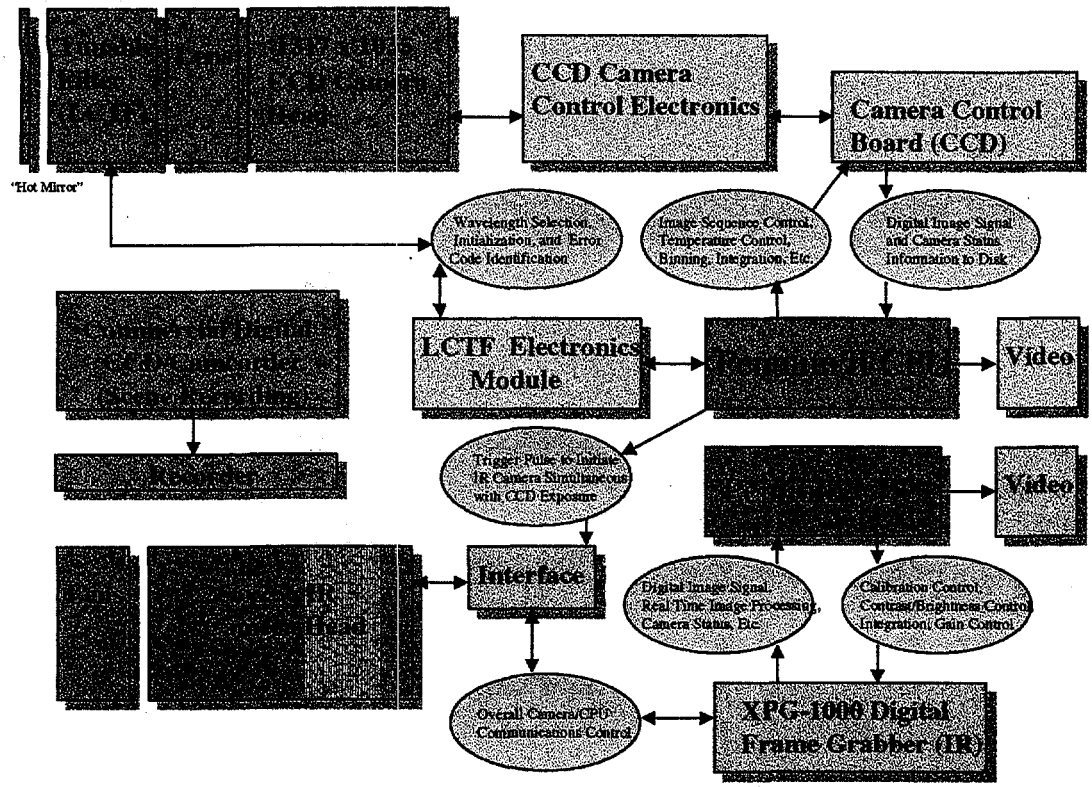

Figure 2.2. Flow Diagram Illustrating the Construction and Data Pathways of the Multispectral Airborne Imaging System. The three camera systems and supporting computers are differentiated by color in the illustration.

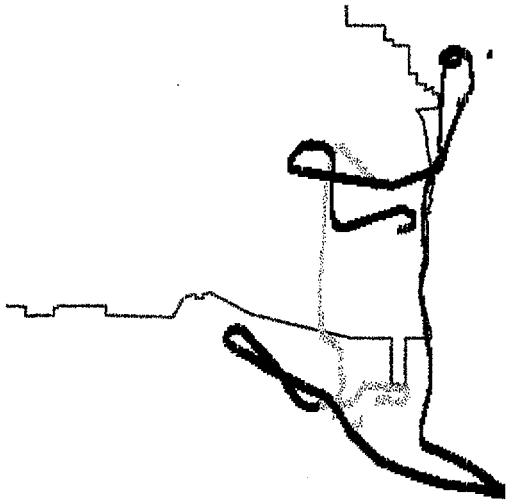

Figure 2.3. Flight Path During a Test Flight Conducted on June 24, 1997. The nearby Hanford boundary is depicted in green; the flight path is shown in red, with flight time increasing from light to dark.

2.2 Software Description

Data collection and camera control for both visible and IR systems are based on Pentium processor computers using the Windows 3.1/95 operating systems. Control scripts written to operate the CCD system rely primarily on the programming language available from within the "V for Windows" imaging processing application software, using Vpascal and PVCAM (Digital Optics Limited and Photometrics, Ltd.). Scripts written for control of the IR system are command-line scripts available from within the "ImageDesk" image analysis application software (Amber Engineering, Inc.). Image acquisition sequences may be operator-controlled or placed in a fully autonomous operation mode. In the latter case, command sequences are issued from the corresponding computers, with timing between the visible and IR systems provided by a single-trigger pulse. During initialization of the imaging system, a number of control parameters specific to each camera (integration period, gain, pixel binning, sub-array readout, etc.) are set, the communications ports are initialized, and the IR camera and LCTF filter are internally calibrated. In addition, a sequence of spectral bands important to the mission is specified and entered into the control scripts; this sequence is continuously looped-through during operations. Synchronization of exposures between the visible and IR cameras is provided by a TTL pulse issued by the visible camera's control computer during each image acquisition command. Image acquisition occurs at a rate of one frame per 1.3 seconds for full CCD resolution and is primarily dictated by A/D conversion times.

After the data are returned to the laboratory, CD-ROM copies of the data are created for analysis and archival purposes. The data may be stored in various formats, depending on the final user's computer platform, preference, and the requirement of the metadata associated with each image.

Future enhancements may include the incorporation of software to automatically select spectral band sequences that are optimized to expected target and background spectral signatures, and automatic selection of integration times to provide for maximum use of the available dynamic ranges.

2.3 Applications to Detection of Proliferation Activities

The objective of the multispectral airborne system is to provide rapid assessments of any changes that might be a result of proliferation activities, both by direct image analysis and through the dynamic clustering and advanced change detection software (described further on in this report). This system is different from existing airborne and satellite multisensor packages in size and ease of deployment. Specific advantages of the multispectral airborne system for Broad Area Search and Analysis include:

- low cost: designed for a light aircraft platform, with minimum staff to operate (2 people; 1 person besides pilot)

- rapid turnaround of data: allows immediate assessment of ground activities and features, rapid redesign of flight missions according to the current information and conditions

- tailored to data needs: tuneable spectral filters in the visible and the availability to collect in the visible, near-IR, and thermal according to the requirements of the desired target (e.g., locating and mapping new utilities and roads in a suspect area)

- spectral bands available may be chosen in real time as conditions and targets dictate

- rapid acquisition of images (i.e., overlap between images) allows directional reflectance information to be exploited for target recognition.

3.0 Landcover and Image Information Extraction

3.1 Classification of Washington Public Power Supply System Using HYDICE Imagery

3.1.1 Background

Hyperspectral systems offer new detection and identification capabilities through tools that match the spectral properties of known targets with image pixels. However, in order to utilize these 'image spectroscopy' capabilities, the imagery has to be calibrated to surface reflectance. The calibration can be through atmospheric (radiative transfer) models or by direct calibration to ground reference targets; the latter approach is more computationally efficient given some knowledge about targets within an image. The focus of this section is to demonstrate hyperspectral image exploitation for landcover analysis, discussing some of the challenges and the use of scene targets for calibration of hyperspectral sensors.

This study takes advantage of a unique dataset of two HYDICE (Hyperspectral Data Image Collection Experiment) flightlines collected during July 1996 over the Washington Public Power Supply System (WPPSS) nuclear reactor located on the Hanford Site. The HYDICE sensor was flown at two altitudes, providing pixel resolutions of 0.75 and 3 meters. The objective of this study was to demonstrate the capability of landcover classification using image spectrometry and its potential usefulness for detection and characterization, especially focusing on the use of existing or derived datasets for image calibration and classification.

3.1.2 Methods

The higher spatial resolution HYDICE imagery (75-cm pixel resolution) was used as a reference image. This image was corrected to reflectance and then classified. Image spectra were then extracted for the calibration panels and key targets; these spectra were then used to classify the higher altitude (3 m) imagery. The methods and results are described in more detail below.

3.1.2.1 Calibrating the Reference Image

On the HYDICE overpass date, ground reference reflectance panels were set up, and ground-based spectral measurements were collected for both the light and dark panels. The HYDICE imagery was later received by Pacific Northwest National Laboratory (PNNL) as digital radiance data (2-byte integer) from the Naval Research Laboratory (NRL). Calibration of the imagery from radiance values to reflectance was performed using a two-point empirical line method, selecting image pixels from the light and dark reference tarps, and relating these to the measured reflectance.

3.1.2.2 Classifying the Reference Image

Several representative scene components were selected as end members, including several building roofs' materials, soil, vegetation, asphalt, and gravel types. An unsupervised classification was used as guidance for selecting the components. Figure 3.1 shows the spectra for these end members; note the similar shape for spectra of man-made features. Spectral Angle Mapping (SAM) was applied to the reference image using the scene end members. The results demonstrate that SAM is insensitive to amplitude differences among spectra so it is not able to distinguish among man-made features with similar spectral shapes. A minimum-distance supervised classification was also used with the same image end members. Minimum distance was chosen over maximum likelihood or Mahalanobis distance because the statistics require (N bands +1) pixels representing each end member (supervised training set); for many of the training sets (such as an individual automobile) this is not possible.

To use the strengths of the two classification approaches, the two classifications were combined on a class-by-class basis. The minimum distance approach was used to separate some man-made features (because this method is sensitive to spectral amplitude differences) and the SAM was used to identify the other features more accurately based on spectral properties. Figure 3.2 shows this resulting combined classification for the reference image.

3.1.2.3 Calibrating the Source Image with Derived Spectra

Although the calibration panels are visible in the source image, given the 3-m size of the pixels, there were few 'pure' pixels comprising the panels; scrutiny of the radiance values indicated that the dark panel was entirely mixed, and the white panel had one unmixed pixel. Rather than use the mixed pixels of the calibration panels, in-scene targets were used instead. A building shadow (of the reactor building) and a

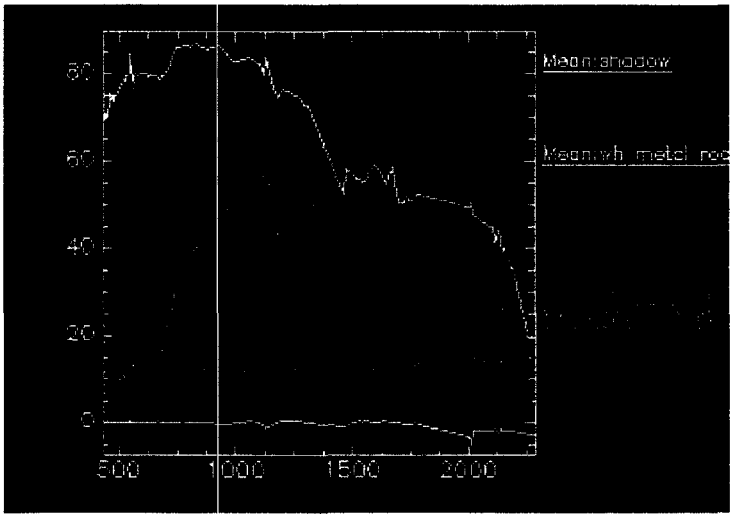

Figure 3.1. Source Image Training Sets Showing the Similarity of Man-Made Features

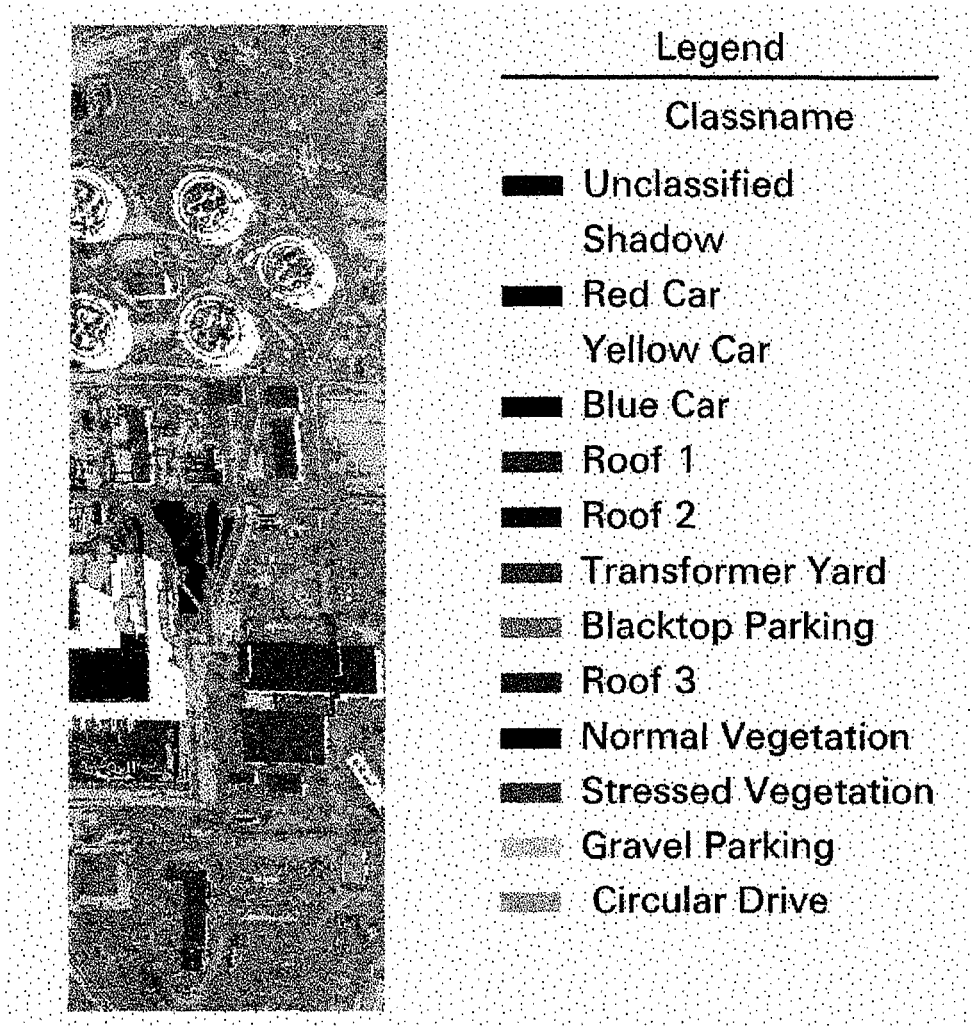

Figure 3.2. Classification Resulting from Combining Spectral Angle Mapping and Minimum Distance Classifications of the Reference (75-cm resolution) Image

metal roof were used as invariant targets. The reflectance spectra for these in-scene calibration targets were derived from the reference image; areas were chosen with more than 100 pixels each. A two-point empirical line calibration was performed to relate reflectance to radiance values for the source image. To check the accuracy of the calibration, the reflectances of the white calibration panel were compared between the source and reference images; this comparison indicated that the panel reflectance in the source image was lower (darker) than in the reference image. However, further analysis of corresponding vegetation areas (much larger than the panel) produced very similar spectra between the source and reference images (see Figure 3.3), so the differences in the white panel spectra were thought to be a result of the mixed pixel problem in the source image.

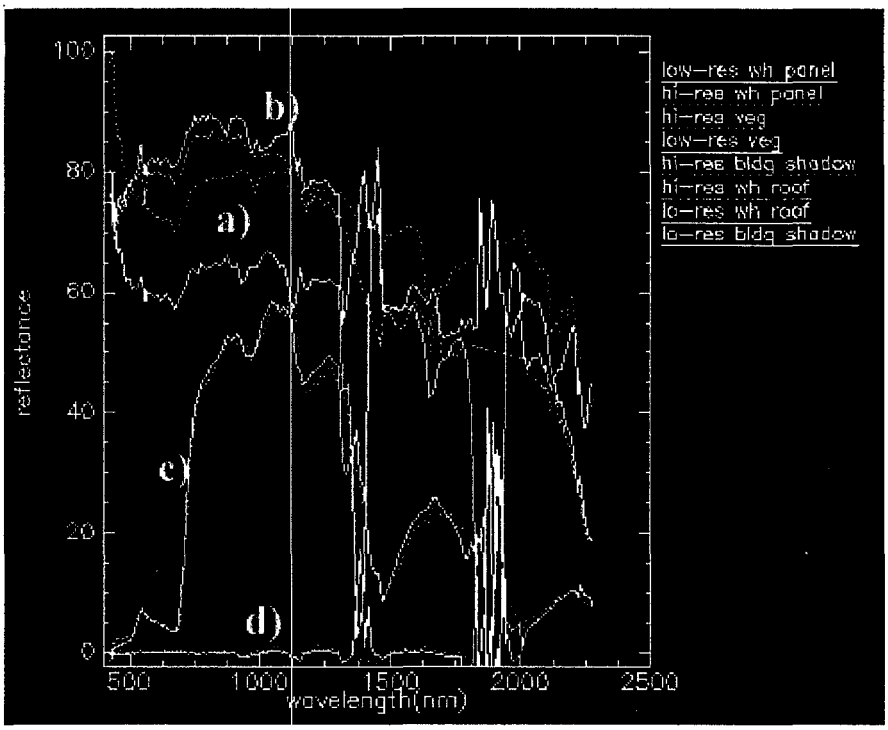

Figure 3.3. Image Spectra of a) White Calibration Panel, b) White Roof, c) Vegetation, and d) Building Shadow Demonstrating the Mixed Pixel Problem Apparent in the White Calibration Panel (note the spread between low- and high-resolution pixels)

3.1.2.4 Classifying the Source Image with Derived Spectra

A set of 12 signatures, representing various targets of interest, was selected from the reference image. Using these signatures, a SAM classification was performed on the source image. The results are shown in Figure 3.4. Most areas of the classification agree with the reference image. Some problems of confusion among man-made features was evident; this would be expected from SAM. The classification did discriminate some cars from the asphalt parking lot.

To test the ability to detect sub-pixel targets, a linear unmixing was performed using the same 12 signatures. The fractional end members for red, blue, and yellow cars indicate that the technique was able to detect some automobiles, despite mixing with the asphalt parking lot (see Figure 3.5).

3.1.3 Conclusions

Spectra from a higher resolution HYDICE image over the WPPSS nuclear plant was used to successfully calibrate and register a coarser resolution HYDICE image of the same area. In-scene targets were used to calibrate the coarser (3-meter) imagery; the reflectance values for these targets were extracted from

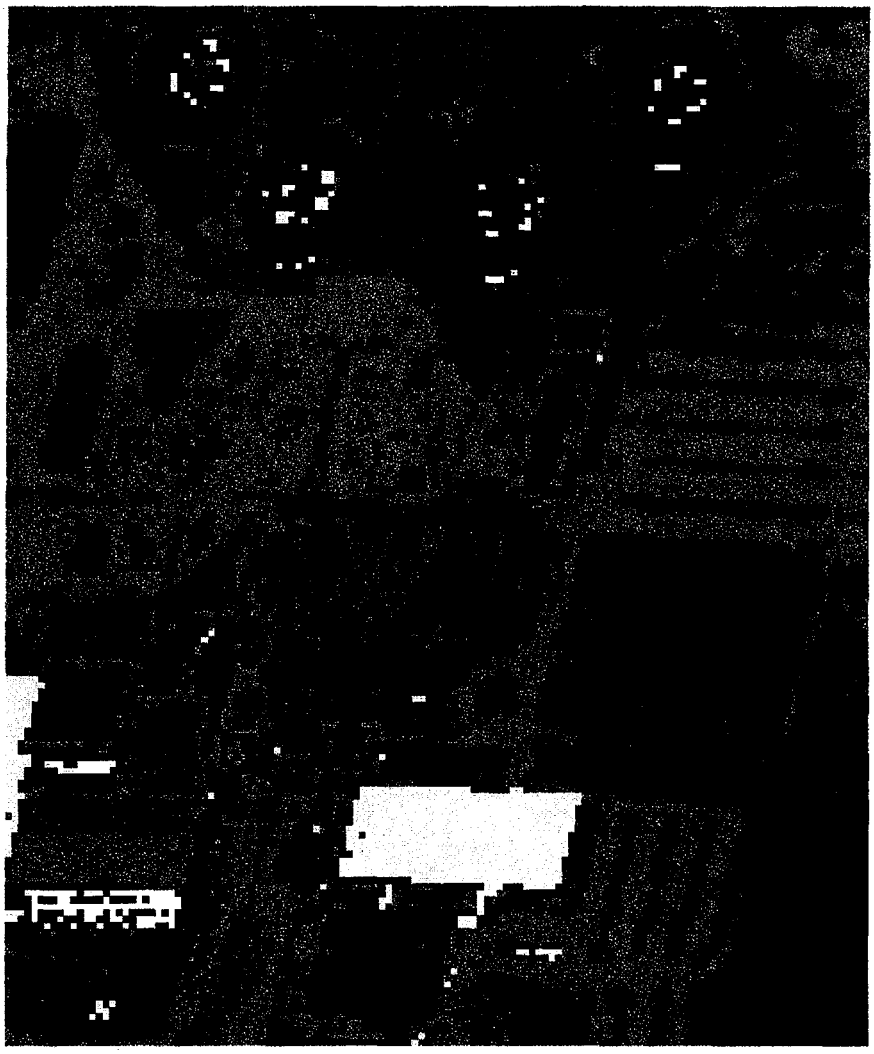

Figure 3.4. SAM Classification for the Source (3-m resolution) Image

the higher resolution (0.75-meter) imagery. Several signatures from the high-spatial resolution image were also used to classify the lower resolution image using both SAM and linear unmixing. Using linear unmixing with the derived end members from the high-resolution image, sub-pixel objects were detected in the coarser resolution imagery. These results have interesting implications for monitoring. One high-spatial and spectral resolution image may be used to calibrate and classify subsequent coarser resolution images, such as those that will be available from upcoming satellite-based sensors.

The WPPSS classification also demonstrated that man-made features often produce spectra that are scalar multiples. In this work, a combination of SAM and minimum distance was found to be useful for separating these similar man-made features while utilizing image spectroscopy was used for separating other features.

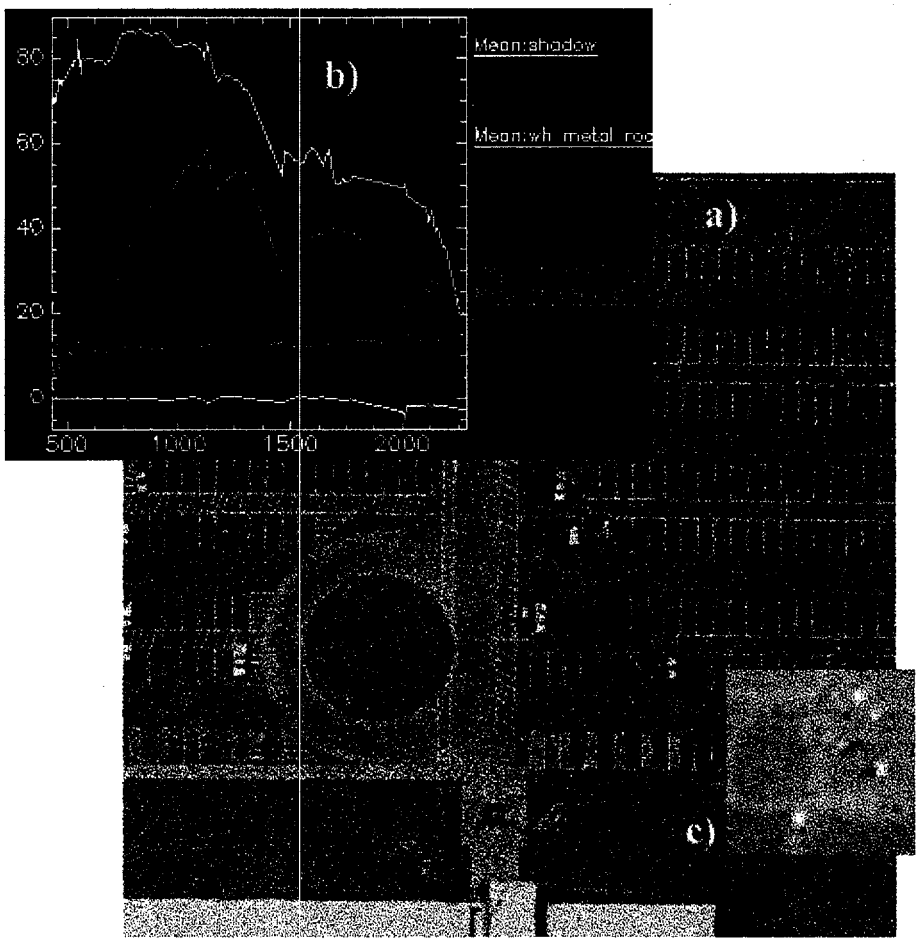

Figure 3.5. a) High-Resolution Photography of a Washington Public Power Supply System Parking Lot Showing Location of Automobiles Near Time of High-Altitude HYDICE Overpass, b) Set of 12 Signatures Used as End Members, and c) Subset of Yellow Car Fraction Image Corresponding to a Portion of the Parking Lot

3.2 Dynamic Clustering: A Rapid Multi-Classifier Approach

3.2.1 Need for a Dynamic Clustering Approach

The primary objective was to develop a near real-time classifier that would generate landcover data sets for use in BASA change detection analyses. This technology was designed to address several needs:

- computational efficiency
- good separation of man-made features
- ability to identify potential image end members
- ability to apply the classifier to a complete image cube or a transformed data set.

Numerous classifications can be used to generate landcover, and each has its own set of problems that impacts image classification. For example, spectral angle mapping (SAM) is a computationally efficient approach that works well with hyperspectral datasets (Kruse et al. 1993); however, one of the unique characteristics of the SAM approach is that it is insensitive to spectral amplitude, that is, two spectra that are multiples will produce the same angle when compared to a pixel to be classed. This is problematic for man-made features (such as asphalt and concrete), which have similar spectral shapes. Other parametric approaches, such as maximum likelihood (Jenson 1986) and Mahalanobis Distance (Swain and Davis 1978), are more computationally intensive for images with tens or hundreds of bands. Also, the statistics are valid only when the number of pixels representing the signature or end member is greater than the number of bands plus 1. This is a problem for many end members that are often determined from as few as one pixel.

3.2.2 Dynamic Clustering Approach

Because no single classification scheme meets the application needs described above, the dynamic clustering approach combines three classifiers: SAM, maximum likelihood, and minimum distance (Jenson 1986) to maximize the classification capabilities. Each of these classifiers is described briefly below:

- SAM determines the angle between the mean in hyperspace (i.e., the mean of each of the N image bands) for an individual pixel and the means (in hyperspace) of each of the image end members. The pixel is then classed to the end member that produces the smallest angle; if the pixel lies outside of a user-specified angle from all of the end members, it is not classed.
- The maximum likelihood classifier determines the likelihood that a pixel belongs to a particular class (represented by the signature or end member); the weighted distance metric that is calculated between a pixel to be classed and each of the signatures is based on both the means and the covariance matrix of the signature pixels.

- The minimum distance classifier computes the euclidean distances between a pixel and the cluster means (in hyperspace) of each of the signatures (or end members).

The dynamic clustering strategy is outlined in Figure 3.6. The process starts with the image to be classified and a set of signatures or end members. These signatures include the mean for each band and the number of pixels that were used to determine the signature. PNNL has developed an autonomous approach for developing image end members (discussed elsewhere in this report). There are also three user-selected parameters:

- *Cluster acceptance angle*, a pixel that lies outside all of the signatures by an angle greater than or equal to this value should be clustered as unclassified

- *Resolution angle*, the minimum angular difference between two signatures that will allow separation by SAM.

- *Resolution length*, the minimum length of a vector (either pixel or cluster) that required for classification by SAM.

Because of inherent noise in the instrument systems and atmosphere, various transforms might be applied to the image. These include:

- Noise reduction, such as principal components analysis or minimum noise fraction (Green et al. 1988)

- Atmospheric corrections (e.g., MODTRAN, ATREM, etc.)

- Correction to ground reflectance values

- Bandset reductions or band selection (Lundeen et al. 1996b)

The image pixel values must be in the same units and waveband regions as the end members so any transformation that alters these image characteristics will be applied to the end members as well. Also, before classification begins, each of the signature vectors are checked to flag the signatures that have only amplitude differences. For each set of two or more similar signatures, a representative signature is selected to use in the classification. In this way, a set of signature vectors is determined for which all of the spectra have unique curves.

After the initial processing (the above steps occur only once for each image), the classification begins. For each pixel, the pixel vector is computed. If the pixel is not within the cluster acceptance angle of any of the signature vectors, it is flagged as unclassified and is added to a set of potential end members. For pixels matched with one of the signature sets, the signature is checked to see if it has been flagged as a multi-signature vector. If not, then that pixel is classed with that signature. For pixels that best match one of the multi-signature vectors, either maximum likelihood or minimum distance classification is used,

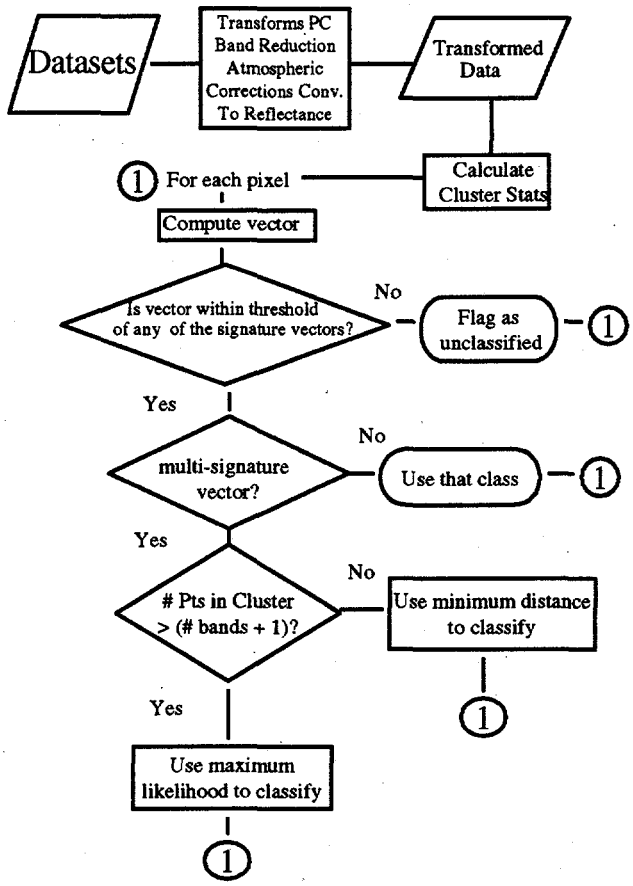

Figure 3.6. Logic Diagram for Dynamic Clustering

depending on how many pixels were used to develop the signature. If all of the signatures in the multi-signature vector were developed from a minimum of (N bands +1) pixels, then maximum likelihood is used, otherwise, minimum distance is used.

3.2.3 Test Case: WPPSS Reactor Facility

To test the proposed approach, PNNL developed software to perform the combined classification. For this test case, only SAM and minimum distance classifiers were used; future implementations will include maximum likelihood classification. That is, if a given pixel fell within the angular tolerance of a cluster vector that represented two of more signatures, then minimum distance was used to classify that pixel. The software was tested on imagery over the WPPSS nuclear power plant at the Hanford Site. The imagery used was 75-cm resolution HYDICE collected during July 1996. For the test case, 12 image spectra were extracted to represent the main features of the area covered by the HYDICE frame.

For the user-supplied parameters, a series of threshold values (0.05, 0.1, 0.2, and 0.4 radians) was used for the cluster acceptance angle. A range of values for resolution angle was also selected (0.01, 0.05, 0.08, and 0.1 radians); the resolution length was set at 10 (the image data had been scaled from 0-100). Figure 3.7 shows six resulting classifications based on the 12 initial clusters and 6 different combinations of the user-selected parameters.

3.2.4 Results and Discussion

The classification results are shown in Figure 3.7. These results show that the selected parameters will influence the classification. As cluster acceptance angle increases (e.g., Figure 3.7a versus 3.7b), the number of unclassified pixels decreases. For this test case, an angle of 0.2 radians (or larger) captures all of the pixels. This may not always be desirable however, because there may be pixels that are not best represented by any of the current clusters. Note the number of unclassified pixels on the left-hand edge of the image (shown in Figure 3.7b and 3.7c), which could be explained by a calibration effect; later this same area appears to have a confusion between transformer yard (gravel) and asphalt roof. As the resolution angle increases, the number of pixels that are classified using SAM decreases, which most affects the man-made features; note the difference between Figure 3.7b and 3.7c. For this classification, the results in Figure 3.7d (cluster acceptance of 0.2 radians and resolution angle of 0.05 radians) agree well with knowledge of the landcover in the area of the reactor.

The initial results from this test show the great potential to develop rapid and accurate surface characterization from hyperspectral imagery. Given signatures of known targets, classification of an image to the desired features (classes) can be performed more accurately in one pass than with other classifiers. For non-proliferation needs, this provides accurate image classifications using hyperspectral datasets in a real-time or near real-time environment, providing up-to-date information on surface activities and conditions. This information is critical in performing change assessments and in establishing the likelihood of undeclared facilities.

Given the encouraging initial results, several improvements to the approach are planned. The first is to include adding the maximum likelihood classifier. Default values for the user-supplied parameters will be provided based on image statistics and the number and characteristics of the features (end members). A direct interface with the end member extraction software (discussed separately in this report) would be useful in that end members could be developed autonomously and incorporated, and the unclassified pixels could be fed back to the end member extraction. Results from the WPPSS test case indicate that there is still some confusion among man-made features; for further discrimination, potential methods based on spatial information would be very useful.

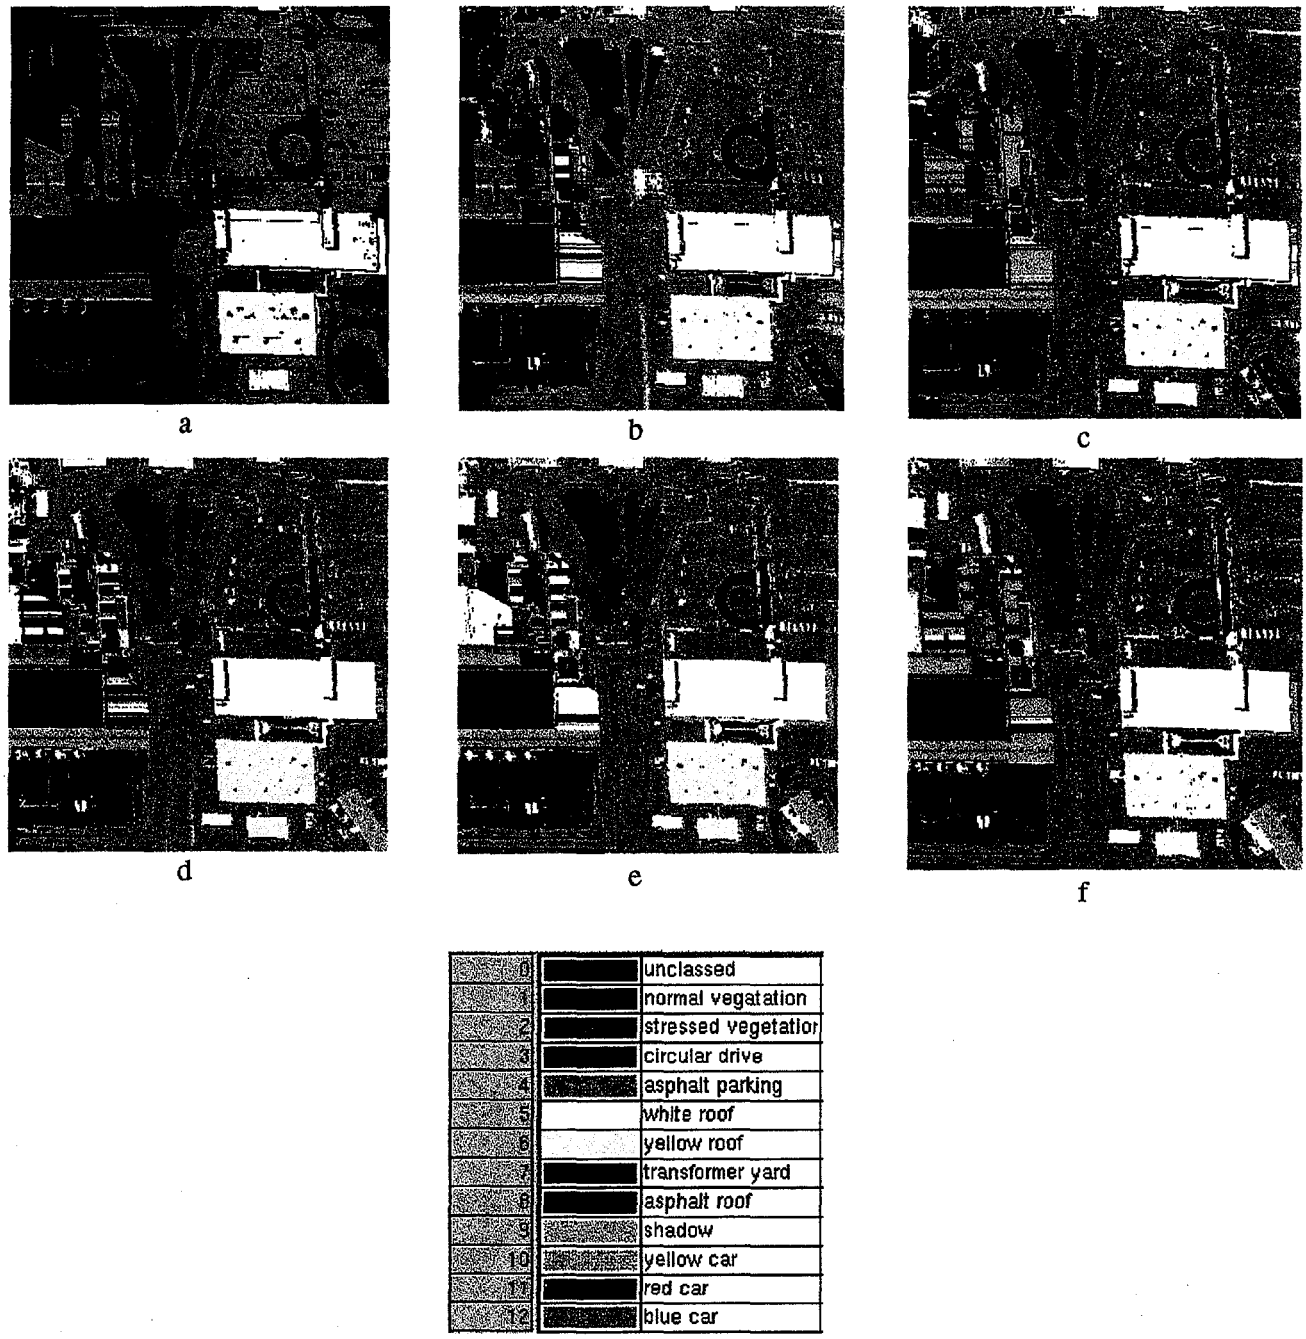

Figure 3.7. Results of Dynamic Clustering for WPPSS HYDICE Imagery. The six frames represent six different runs of the classifier using different values for the cluster acceptance (the first number) and resolution angle (the second number) parameters. These values are 2a) 0.05, 0.01; 2b) 0.1, 0.01; 2c) 0.1, 0.05; 2d) 0.2, 0.05; 2e) 0.4, 0.08; 2f) 0.2, 0.1. Note that as the cluster acceptance angle increases, the number of unclassified pixels decreases; as the resolution angle increases, the number of pixels which are classified using SAM decreases.

3.3 Hyperspectral Image Dimensionality Reduction with End Member Analysis

3.3.1 Introduction

The greatly increased spectral resolution of hyperspectral data sets offers the promise of greatly improving identification and characterizing important features. For example, the HYDICE sensor divides the light energy from the visible to the near IR into 210 bands (Figure 3.8). With this fine resolution it is theoretically possible to resolve features by very fine differences in 'color'. However, such improved capabilities do not come without cost. The large amounts of data that must be processed present issues of: increased operational complexity and cost, computational burdens, bandwidth data transmission limitations, and loss of capabilities that might be offered though design trade-offs. For instance, there is often a trade-off between the number of bands that are collected and the spatial resolution. Further, while higher spectral resolution may permit delineation of narrow spectral features, this advantage often comes at the cost of reduced signal-to-noise ratio compared to broader band measurements. These problems have provided strong motivation for PNNL to investigate ways to exploit the correlation structure in the data so that the number of bands that have to be collected to characterize a given problem can be reduced to 20 or fewer and still retain the ability to discriminate important features (Lundeen et al. 1996b). This work successfully investigated methods on how to best exploit known information to minimize the collection of bands that contain redundant data and maximize the collection of bands that have the greatest discriminating power. This section describes work that builds on this experience to investigate another approach based on the idea of defining a hyperspectral image in terms of low dimensionality end members.

3.3.2 End Member Strategy

With the end-member strategy (e.g., Smith et al. 1985, 1987; Pech et al. 1986), each pixel in an image is considered a mixture of different cover types or spectral end members. Specifically, in linear mixing the value of each pixel in an image is a linear combination of the end members, and the weight given each end member is directly proportional to the area covered in the pixel by the corresponding cover type. For hypothetical example, consider a two-band sensor system used to capture an image of a farming district with ground cover consisting of only trees, wheat, and corn. If a pixel in the image covered only half of a corn field and half of a wheat field, then the values of the bands collected in the corresponding pixel would be

band1 = (0.5* pure corn value of band 1) + (0.5* pure wheat values of band 1)
band2 = (0.5* pure corn value of band 2) + (0.5* pure wheat values of band 2)

In Figure 3.9 the possible pixel values are given by the triangle formed by the pure end members. With standard linear unmixing techniques it is possible to transform the original two-band image into three bands

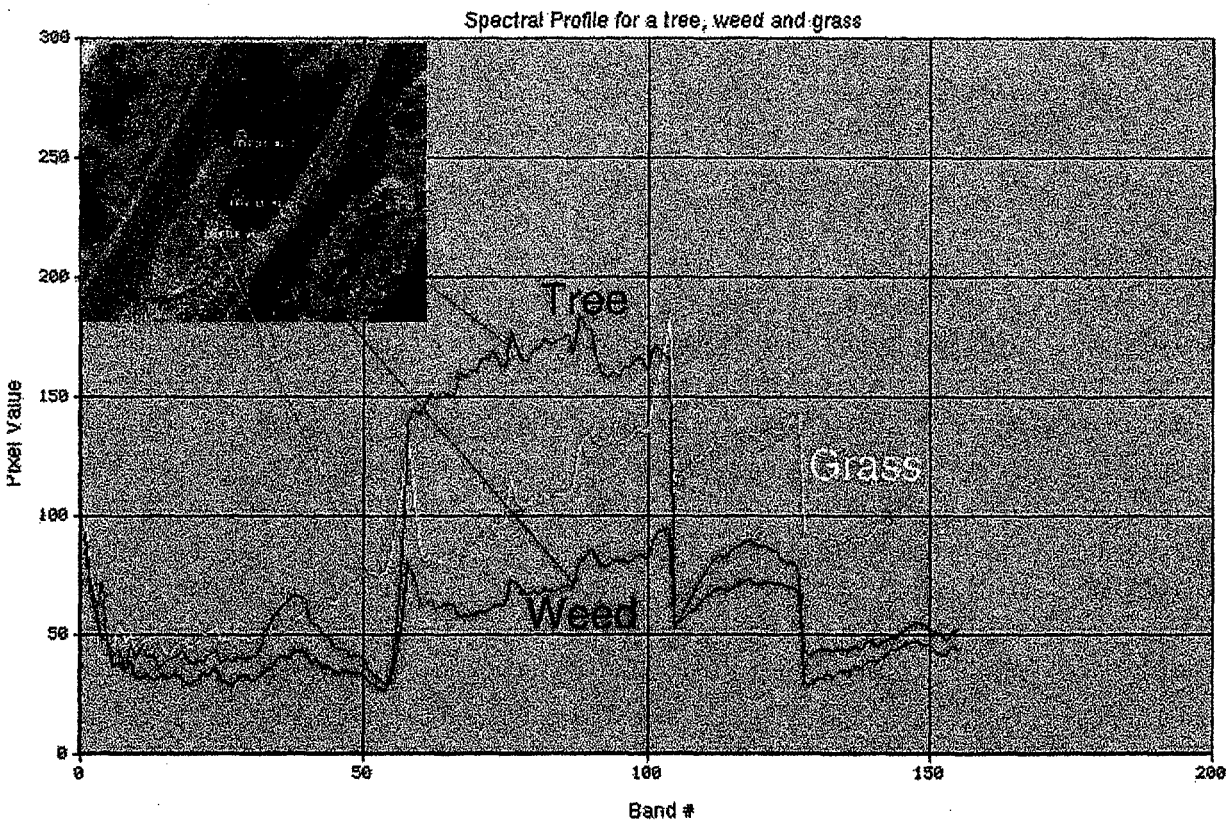

Figure 3.8. Example Spectra from the HYDICE Sensor (fewer bands affected by water vapor absorption)

that reflect the percentage of corn, wheat, and trees in each pixel. This simple example represents the optimum case in which two independent bands are transformed into new bands with direct physical meaning. Since, in this example, the bands are not correlated (i.e., the dimensionality of the system is really two) it is possible to transform the collected data exactly into three new bands (i.e., the number of bands plus one). However, with hyperspectral data the bands are often highly correlated. In some sense there is a great deal of redundant data, and the true dimensionality of the image is considerably less than the number of bands. Previous work by PNNL suggests that this redundancy implies that less than 20 end members can often adequately represent a scene (Lundeen et al. 1996b). Thus, once the correct number of end members have been properly identified, it is possible to reduce the dimensionality of the problem to the number of the end members. Further, the new bands have a direct physical interpretation (e.g., a band where each pixel gives the percentage of wheat). However, in applying this strategy several problems typically occur. For instance, as with a supervised classification, the user must specify the end members. Our experience strongly suggests that if the end members are incorrectly chosen, or characterized, then the results can be problematic. In a complex scene the selection of the correct number and type of end member ground cover spectrums can be difficult. Further, in some cases the problem set only requires the identification of one target, and the effort to properly characterize other end members is inefficient.

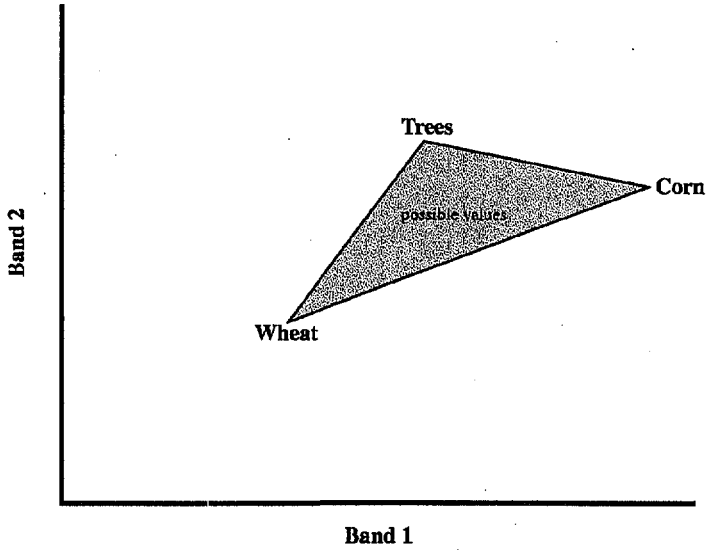

Figure 3.9. Plots of Pure Corn, Wheat, and Tree Spectral for Pure Corn, Wheat, and Trees for an Idealized Two-Band System

To address these problems, an experiment was investigated with a way to automatically pick end members building on a hyperspectral image classification and dimensionality reduction method proposed by Harsanyi and Chang (1994).

3.3.3 Orthogonal Subspace Projection

The Harsanyi and Chang (1994) methodology starts with the assumption that a user can provide the end members, one of which is of direct interest (Figure 3.10). For this experiment, a HYDICE hyperspectral image was used with seven end members: trees, grass, water, a dirt road, a paved road, shade, and Scotch broom. This example was chosen because:

- vegetation types are notoriously difficult to separate, even using spectral methods
- changes in vegetation can be a key indicator of man-induced changes, such as the presence of toxic chemicals, underground structures, or increased off-road vehicle traffic
- a DoD client needs timely and accurate assessments of noxious weeds for resource management.

In this example, the goal is to find the amount of these weeds over the site in order to plan and budget weed control activities. Given the true end members, the method first projects the spectral data of each pixel into a subspace that is orthogonal to the undesired signatures. In our example, these include the spectral signatures for trees, grass, water, dirt roads, shade, and paved roads. Once the interfering signatures have been nulled, the residual is projected onto the signature of interest (i.e., weed spectra). This results in a single component image that best represents (in a signal-to-noise ratio sense) the ground cover of interest

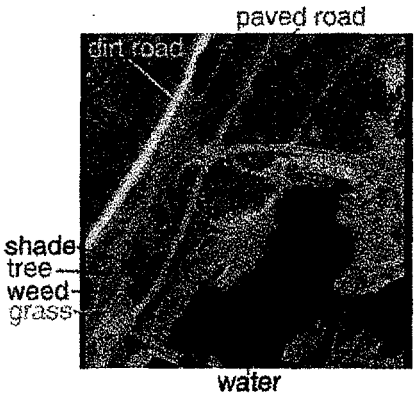

Figure 3.10. False Color Image of Example Study Site

(i.e., weed content in our example). However, the methodology depends on the user supplying a complete and accurate description of the end members. Leaving out end members can lead to misleading results. For instance, in testing the methodology with the example image (Figure 3.10) the tree end member was accidentally left out in one case. The resulting image was not a satisfactory representation of weed content in that the tree pixels, which should have been close to zero in weed content had high values in the weed bands. This error was corrected when the tree spectral was added to the end member list. However, this problem case also suggested a mechanism for selecting end members. Specifically, those pixels that were misidentified as having high weed content could be used to identify candidates for empirically based end members.

3.3.4 Candidate Method for End Member Selection

In extending the Harsanyi and Chang (1994) method to select end members, the first step is to provide an adequate spectral representation of the main feature of interest (Scotch broom, in our example) and define areas in the image where pixels have been classified either as weed or candidate pixels (Figure 3.11). The procedure then iterativly follows the next steps:

- Calculate the single band representation for the target (Scotch broom in this example)
- Find the 100 largest values in the background candidate pixels (which would optimally be zero if the end members were specified correctly). Of these high-error pixels, choose the pixel that is most unlike the target (Scotch broom) as the new end member; use the spectral angle difference to choose the pixel that is the most unlike the target
- Continue the procedure until either the change is small or the image is reduced to noise.

Figure 3.11 shows the results of the first step. Figure 3.12 show the results after nine iterations. Note that the definition of Scotch broom is significantly improved.

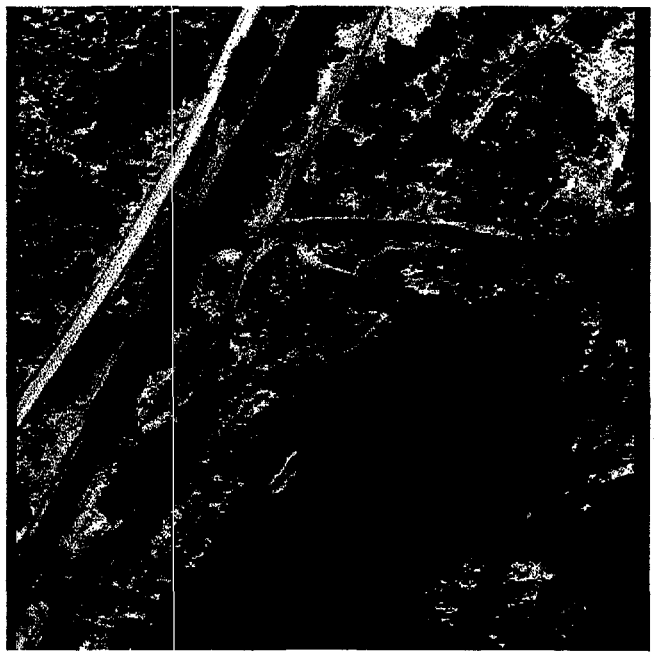

Figure 3.11. Weed Content Color Coded from Low (blue) to High (red) with End Members Missing

Figure 3.12. Weed Content Color Coded from Low (blue) to High (red) Calculated with all End Members

3.3.5 Conclusions

The standard Harsanyi and Chang (1994) methodology is analogous to the supervised classification methodology, in that it requires the user to provide spectral characterization of each class. This extension is more analogous to an unsupervised classification since it iterativly selects its own best end members in some mathematical sense. Like unsupervised classification, it also does not provide a physical interpretation of the spectral end members. However, the process of attempting to attach physical meaning to the derived points may be of value in its own right.

To date, PNNL has generated prototype software and experimented with a single scene. This effort has been extremely promising. However, to fully test and develop the methodology more work is needed, preferably using an example of a proliferation facility.

4.0 Change Detection Optimized for Broad Area Search

4.1 Introduction

An important part of the U.S. Department of Energy's (DOE's) mission is the identification and characterizing of nonproliferation activities around the world. This challenging effort includes the monitoring of board areas of the earth's surface, with diverse physical and cultural characteristics, for suspicious activities. This inherently difficult effort is further complicated by budget, human resources, and time restrictions. One important tool used to meet these challenges is change detection using remotely sensed images. Examples include the use of both civilian SPOT imagery and NTM data sets. However, these data sets have limitations. For instance, SPOT is limited by relatively coarse 10-meter pixels, foreign control, significant costs, scheduling issues, high image registration costs and limited revisit times. NTM data sources avoid some of these issues but the extensive use of NTM data sources is severely constrained by other factors. However, with the expected launch of a suite of new satellite systems in the near future (see our web site at http://www.pnl.gov /remote/projects /paper/begin.htm for more details) many more capabilities will be available to meet DOE's mission requirements. Examples of increased capabilities include improved

- spatial resolution (1-meter pixels)
- ground location accuracy
- frequent revisit times
- global coverage
- band options
- stereo
- signal-to-noise ratios and dynamic range.

While exploiting these new capabilities will greatly increase opportunities to monitor events around the world, they will also present DOE with a significant challenge to effectively exploit the 1) changing nature of the bands collected (e.g., hyperspectral data sets), 2) vast amounts of data (e.g., 100 times the amount of data in a SPOT image), and 3) integration with other new spatial data sets. In particular, several ongoing trends support the exploitation of these new satellite systems. These trends include

- cost reductions and more powerful computer hardware

- increased growth in geographic information system (GIS) and global positioning satellite (GPS) technologies
- competitive forces that promote better turnaround times and reduced cost
- mature software
- transition from analog to digital data collection
- pool of experienced personal
- availability of supporting digital data sets (e.g., DEM).

Of particular importance are GIS data sets, which are growing in use and availability around the world. DOE must be in a position to fully exploit these data sets. This section describes a change detection methodology that is being developed to address these issues to take advantage of new GIS capabilities.

4.2 Change Detection Methodology

The overall methodology for improved change detection is shown in Figure 4.1. The first step is to take two registered images (the step called raw difference in Figure 4.1) and create a first approximation image of the differences (Lundeen et al. 1996a). In this step, several options are available, including 1) subtraction of images, 2) image ratios, or 3) line fitting. In line fitting, the two images are used to generate a line of points by plotting the corresponding pixel pair in a graph (i.e., let a pixel from time 1 give the value for the x axis, and let the y value be taken from the pixel at the same location from the image at time 2). If the images were exactly the same, then the line would be a straight line though the origin at 45 degrees. If only broad atmospheric changes occurred between images, the offset and slope of the line would change. Local changes (e.g., new roads) would displace local pixel pairs away from the broad trend. Thus, to identify local changes and minimize the effect of broad-scale atmospheric changes between images, the line fitting procedure 1) plots the data, 2) fits a line (which normally accounts for atmospheric differences), and 3) creates a new image in which each pixel value represents its distance from the main trend line. This new image can then be level-sliced to identify changes. The main advantage of this line fitting methodology is that is tends to minimize the changes associated with broad atmospheric effects and enhance the effects of local changes (e.g., new buildings) that are of more interest. Figure 4.1a illustrates the procedure with an example using the 200 Area nuclear site from Hanford. One way to improve on this line fitting methodology is to treat the change detection process as a classification problem. In the classification change detection methodology (the classification process in Figure 4.1) a scatter diagram is again made with the two images. However, using the first raw difference image (Figure 4.1) has a guide, the user interactively outlines in the scatter diagram areas that represent changes. Typically, four classes are identified: major

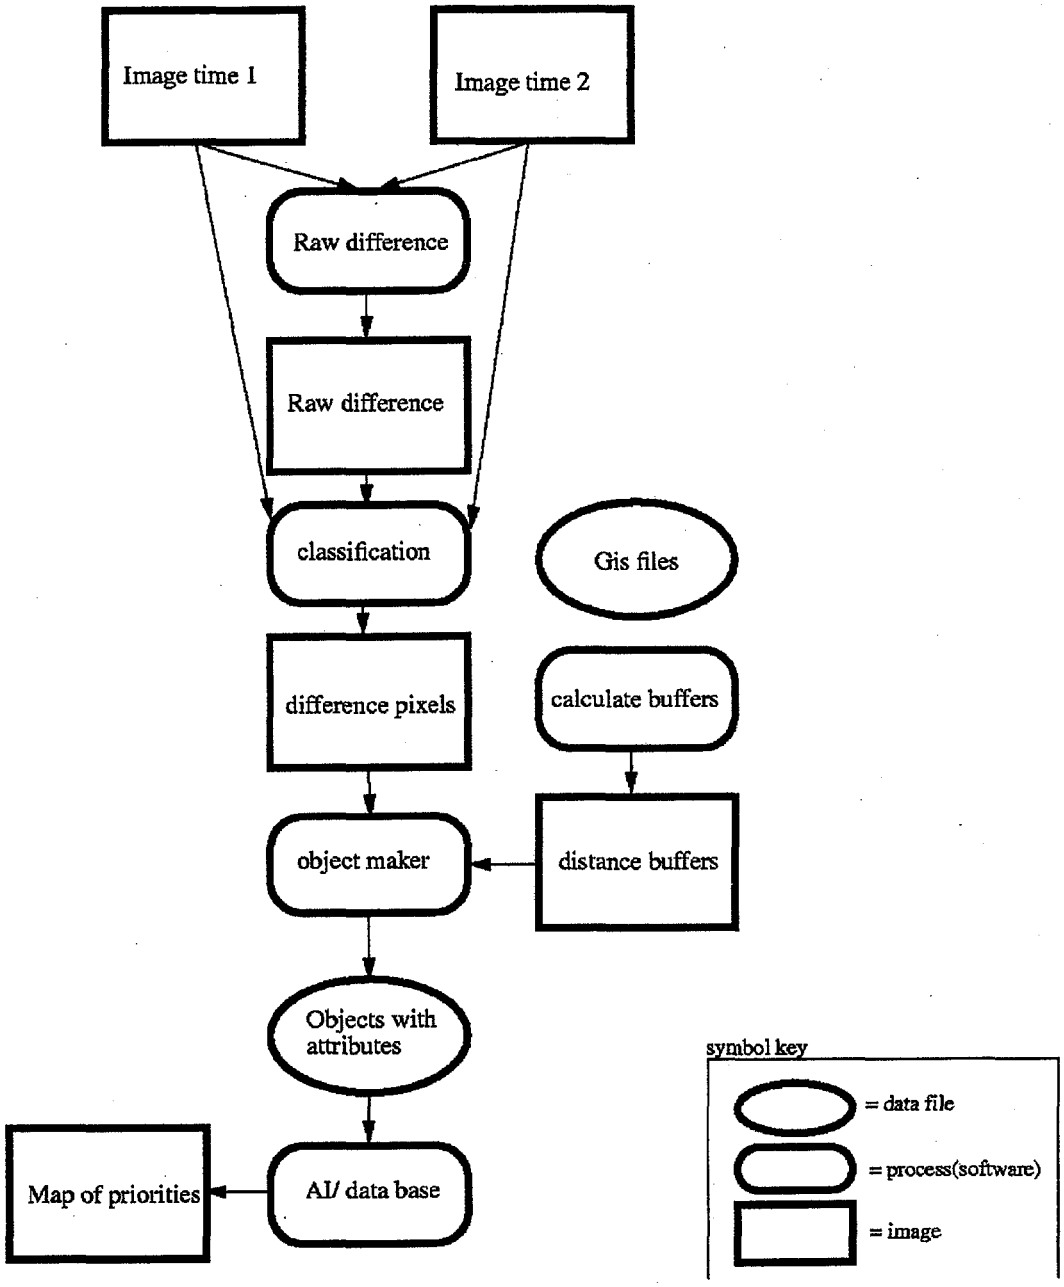

Figure 4.1. Change Detection Data Flow

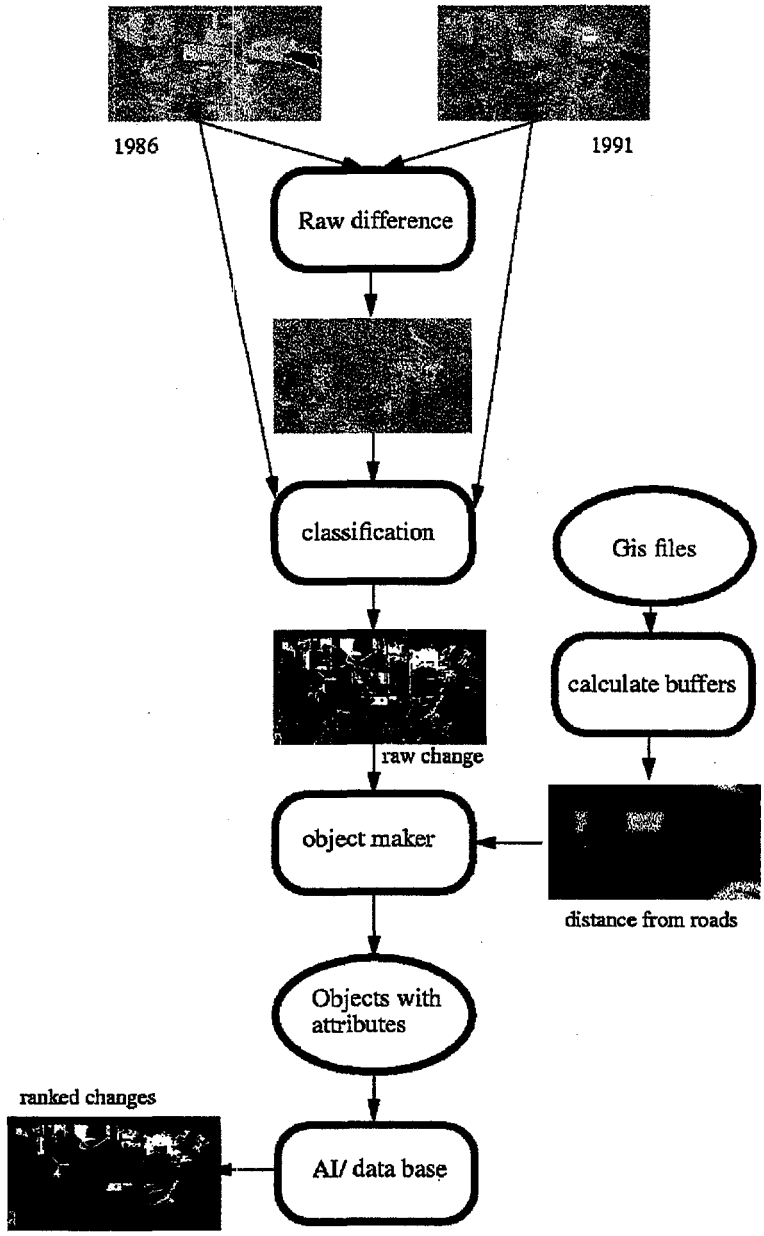

Figure 4.1a. Change Detection Data Flow with Example Images

increases in energy (e.g., a new white roof on a building), minor increases, minor decreases and major decreases, (e.g., a new asphalt parking lot). This strategy allows the user to define changes much more finely, but at increased labor cost.

Once a raw difference map is created, the next step is to group all the individual pixels into objects. This is done by grouping all individual pixel that touch each other. Two options define touching. In one case, pixels are defined as touching if they share a common side (e.g., pixels directly to the east, west, north or south). An alternative is to also include pixels that touch at the corners (e.g., northwest, northeast, southeast, southwest). In general, our experience suggests that the simpler option of grouping only pixels sharing sides works best. Regardless of which grouping option is used, the next operation is to characterize each object. The software calculates for each object a number of geometric attributes that include

- area
- surface area
- center of gravity
- moment of inertia
- bounding box.

At this stage, GIS data are also included. In the nuclear site example, distances from major roads and buildings were calculated (buffers image in Figure 4.1a). With this information, the attributes for each object can be enhanced with GIS information that can include

- minimum distance to a road
- maximum distance to a road
- average distance to a road
- minimum distance to a building
- maximum distance to a building
- average distance to a building.

The distance buffer images are normally created only infrequently (the blue path in Figure 4.1), which has an important practical implication: adding GIS attribute information to the objects is relatively easy and inexpensive once the initial investment is made in creating the buffer layers.

Once the images have be transformed into abstract objects, it is practical to perform rule-based selection on them. In our nuclear site example, each object is evaluated with a simple C program, and points are given for each criteria met. For example, if a change is large and near a major building it is ranked higher (i.e., given more points) than a small, isolated change. The final product is an image that shows the ranking of changes between the two scenes (Figure 4.1a). These rankings can in turn be used to optimize the expensive human analysis resources.

4.3 Conclusions

To date, a change detection methodology has been developed that is optimized for exploiting new trends in remote sensing and GIS in broad area search. The end result is designed to optimize the expensive and limited human analytical resources. Current experiments to date indicate the effectiveness of this approach for detecting important changes in the environment; however, the need for more work is also indicated. One example that will be explored in the near future is the use of a sophisticated database engine to enhance the last step of rule-based ranking.

5.0 Conclusions

This report described three main components of an integrated remote surveillance system: an airborne multispectral imaging system, near real-time image information extraction, and advanced change detection technology. The technologies provide important tools that will improve identification of specific targets (e.g., evidence of underground structures) as well as the assessment of areas for suspect activities through inference (e.g., the type of changes combined with intelligence and a priori information).

Future work on this system may include improvements in both data collection and analysis. For the airborne imaging system, the use of GPS coordinates in the PNNL-developed autoregistration software would provide near-autonomous image registration. In addition to complete landcover classification, the extraction of 2-D features (such as roads, buildings, or water features) could be useful for the change detection as another layer of information to compare against benchmark conditions. The DOE CERBIUS project will provide an ideal real-world opportunity to demonstrate the integration of these technologies.

6.0 References

Green, A. A., M. Berman, P. Switzer, and M. D. Craig. 1988. "A Transformation for Ordering Multispectral Data in Terms of Image Quality with Implications for Noise Removal." *IEEE Trans. On Geosciences and Remote Sensing* 26:65-74.

Harsanyi, J. C., and C-I Chang. 1994. "Hyperspectral Image Classification and Dimensionality Reduction: An Orthogonal Subspace Projection Approach." *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 32, No. 4.

Jenson, J. R. 1986. *Chapter 8 Supervised Classification in Introductory Digital Image Processing.* Prentice-Hall, New Jersey. 370 pp.

Kruse, F. A., A. B. Letkoff, J. B. Boardman, K. B. Heidebrecht, A. T. Shapiro, P. J. Barloom, and A.F.H. Goetz. 1993. "The Spectral Image Processing System (SIPS) – Interactive Visualization and Analysis of Imaging Spectrometer Data. *Remote Sensing of Environ.* 44:145-163.

Lundeen, T. F., A. K. Andrews, E. M. Perry, M. V. Whyatt, and K. L. Steinmaus. 1996a. *Development of Automated Image Co-Registration Techniques: Part II – Multisensor Imagery.* PNNL-11383, Pacific Northwest National Laboratory, Richland, Washington.

Lundeen, T. F., P. G. Heasler, and G. M. Petrie. 1996b. *Automatic Band Selection for Sensor Optimization.* PNNL-11360, Pacific Northwest National Laboratory, Richland, Washington. 24 pp.

Pech, R. P., A. W. Davies, R. R. Lamacraft, and R. D. Tgraetz. 1986. "Calibration of Landsat Data for Sparsely Vegetated Semi-Arid Rangelands." *International Journal of Remote Sensing* 7,1729-1750.

Smith, M. O., P. E. Johnson, and J. B. Adams. 1985. "Quantitative Determination of Mineral Types and Abundances from Reflectance Spectra Using Principal Component Analysis." *Journal of Geophysical Research* 90, 792-804.

Smith, M. O., D. A. Roberts, H. M. Shipman, J. B. Adams, S. C. Willis, and A. R. Gillespie. 1987. "Calibrating AIS Images Using the Surface as a Reference." In: *Proceedings, 3rd Airborne Imaging Spectrometer Data Analysis Workshop*, JPL Publications, Pasadena, California.

Swain, P. H., and S. M. Davis. 1978. *Remote Sensing: The Quantitative Approach.* McGraw Hill Book Company, New York.

Distribution

No. of Copies		
ONSITE		
25	**Pacific Northwest National Laboratory**	
	H. P. Foote	K9-55
	J. L. Fuller	K6-48
	C. S. Kimball	K9-55
	R. R. Kirkham	K9-33
	B. D. Moon	K9-55
	E. M. Perry	K9-55
	G. M. Petrie (8)	K9-55
	B. A. Roberts	K6-48
	D. W. Slater	K9-55
	K. L. Steinmaus	K9-55
	M. V. Wyatt	K7-28
	Information Release (7)	K1-06