



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Overhead Detection of Underground Nuclear Explosions by Multi-Spectral and Infrared Imaging

J. R. Henderson, M. O. Smith, M. E. Zelinski

October 14, 2010

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed 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 Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Overhead Detection of Underground Nuclear Explosions by Multi-Spectral and Infrared Imaging

23 August 2010

John R. Henderson, Milton O. Smith, and Michael E. Zelinski, LLNL

Executive Summary

The Comprehensive Test Ban Treaty (CTBT) permits Multi-Spectral and InfraRed Imaging (MSIR) to be performed as part of an On-Site Inspection (OSI) for the purpose of reducing the search area for the location of a possible underground nuclear explosion (UNE). Dedicated airborne MSIR measurements have not been made in conjunction with historical or recent UNE's, so satellite data has been used to determine if there are MSIR observables associated with recent UNE's. In this work MSIR data from commercial satellites has been used to show that there are detectable surface observables which can be used to greatly reduce the search area for the location of the UNE. This has been demonstrated using Landsat data of the Indian, Pakistani, and North Korean UNE's in the last 12 years, and with GeoEye-1 data for the North Korean tests. The techniques used typically identify a region of interest less than 1 km² in size (compared to the nominal 1,000 km² search area), and the few false positives have been resolvable as such by using visible imagery.

The results of this study show that MSIR data from satellites can be used to help prepare the inspection team for an OSI. The Landsat data used here were chosen for their ready availability, the expectation that the satellite spatial resolution and spectral bands would be useful, and the 16-day site revisit time of the satellite. Data from other satellites may have greater utility – for one site, GeoEye data with 3 meter spatial resolution were used to find regions that were not detected in the 30 meter spatial resolution Landsat data. The analysis techniques used here are fairly straightforward change detection or surface categorization techniques and a more sophisticated spatial/spectral algorithm that uses both properties to find anomalous regions in a data set. Improvements in sensitivity and reduction of false alarms are expected with the development of more sophisticated techniques.

It is anticipated that analysis of data from other satellites might be used with the current results to help with the specification of MSIR equipment to be used for an OSI by the inspection team. The GeoEye data show that more sensitive detection of regions of interest is possible with higher spatial resolution, although that data suffers from a minimal set (4) of spectral channels. Additionally, this work has shown that automated algorithms are particularly useful in sifting through large data volumes and detecting specific types of anomalies that are not readily apparent to visual inspection of the data.

1.0 Background and Purpose

The Comprehensive Test Ban Treaty (CTBT) allows for Multi-Spectral and Infrared Imaging (MSIR) measurements as part of an On-Site Inspection (OSI). The objective of MSIR measurements is to help narrow the search area for the location of a potential nuclear explosion that violates the CTBT. The application and utility of MSIR measurements for OSI require further study: MSIR measurements have not been made or used to date for field exercises conducted by the CTBTO; there are few measurements with demonstrated relevance to this objective; there are no vetted requirements for the specification of equipment that might be used for MSIR; and the current level of experience is insufficient to generate CONOPS and analysis procedures for the use of MSIR to support an OSI.

In general, the UNE observables that might be detected by MSIR fall into five categories:

- 1) disturbed earth at the surface (due to the shock wave from the explosion),
- 2) plant stress in the vicinity of the UNE,
- 3) artifacts of human activity,
- 4) thermal effects, and
- 5) novel materials at the surface.

The first has not been explicitly measured for a UNE, but visible observations of earth movement, surface fissures, and measurements of surface upheaval from prior UNE's indicate that disturbed earth is possible for a UNE of sufficient size. Further, measurements for other purposes have shown that disturbed earth can be detected with MSIR imaging.

Airborne MSIR measurements of plant stress were made during the Non-Proliferation Experiment (NPE), so this observable has been demonstrated under relevant conditions. The amount of plant stress was observed to peak one to two days after the explosion, and then fade back to pre-explosion levels after about a week. [William L. Pickles, "Observations of Temporary Plant Stress Induced by the Surface Shock of a 1-kt Underground Chemical Explosion," UCRL-ID-122557, December 1995.]

Observables associated with activities in support of a UNE are expected to vary with the specific conditions of executing the UNE, but MSIR measurements are often able to detect activities such as recent traffic on dirt and gravel roads, as well as thermal indications of activity in buildings. Temperature related observables lie in the long-wave infrared, while road traffic and disturbed earth can have observables in the visible and infrared spectral regions.

Thermal effects at the surface might be due to either hot gases from the UNE blast escaping to the surface, or underground water flows being redirected to the surface by changed sub-surface geology. Novel materials could be brought to the surface through either venting of materials

from the UNE, or by migration of native sub-surface material to the surface because of the escape of hot gases or water.

Maturing the use of MSIR for an OSI requires that the MSIR observables of an UNE be well characterized so the MSIR equipment, deployment CONOPS, and data analysis techniques can be specified. Given the absence of relevant airborne data (except for plant stress), satellite data was used to determine whether MSIR observables are reliably associated with a UNE, and to characterize those observables. Since MSIR instruments on satellites have improved significantly over time, this work has focused on UNE's conducted in the last 12 years.

The purpose of this paper is to demonstrate that (1) there are MSIR observables associated with UNE's, and (2) those observables have unique spectral/spatial properties which can be exploited with a basic set of analysis tools to greatly reduce the search area for purposes of an OSI. A later paper will provide technical detail on the algorithms and techniques used to generate the regions of interest from the satellite data.

2.0 Selection of Data

The six UNE's in the last twelve years were chosen as the focus of this study because they represent the size of UNE (approximately one kiloton of yield) of concern to the CTBT, and because they are recent enough that a variety of potentially relevant satellite measurements exist. While Landsat 5 (whose Thematic Mapper data was used here) was launched in 1984 (well before the cessation of nuclear testing in 1992), satellites with higher spatial resolution and better spectral capabilities have launched more recently. The Landsat data was used because the data is readily downloadable; it has relevant spectral bins, including a standard plant stress data product; the spatial resolution is acceptable (30 m versus an expected feature size of over 150 m for disturbed earth or plant stress), the data volumes are manageable, and the frequency of observation is acceptable for the disturbed earth measurements. Google Earth imagery, and higher spatial resolution data from additional satellites were used to understand the context of results obtained from the Landsat data.

The six UNE's are, in order of ease of detection of MSIR observables (not chronological order):

- 1) 28 May 1998, Pakistan (seismic magnitude $m_b = 4.8$)
- 2) 30 May 1998, Pakistan ($m_b = 4.6$)
- 3) 11 May 1998, India ($m_b = 5.2$)
- 4) 13 May 1998, India (no seismic signal detected)
- 5) 9 October 2006, North Korea ($m_b = 4.1$)
- 6) 25 May 2009, North Korea ($m_b = 4.5$)

The Pakistani and Indian test magnitudes are from W. R. Walter et al, Preliminary Regional Seismic Analysis of Nuclear Explosions and Earthquakes in Southwest Asia, LLNL, UCRL-JC-130745, July 1998. North Korean test magnitudes are from a 25 May 2009 CTBTO press release.

3.0 UNE Detection Algorithms

The techniques used here for detecting observables associated with a UNE fall into three categories. The first is change detection, the second is surface categorization, and the third uses the spatial and spectral properties of the data to search for anomalies. In change detection, reference data is used from before (or after) the seismic event to establish the level and variability in each spectral channel for each pixel in the scene. The post-event data is then compared to the pre-event baseline to determine if any statistically significant changes have occurred for any pixels. These regions are flagged (by their spatial location and statistical significance) for consideration as candidate locations for the UNE.

The second technique is to use the spectral properties of scene to group pixels into a few categories, where the pixels in each category have a similar spectral shape. Categories typically correspond to scene features such as soil, rock, vegetation or shadows. The pre-event and post-event categories are compared spatially, and any differences evaluated. For example, a region that showed as vegetative pre-event and then shows as soil post-event would be a candidate for the UNE location since surface disruption might overturn vegetation and expose bare ground. An advantage of this technique is that it does not require the use of pre-event data, although it is probably more sensitive when pre-event data is available. For example, a region that shows up as soil might be a candidate for proximity to the UNE if it can be distinguished from roads or farms by its size, shape, location or spectral difference from adjoining regions.

The third technique is similar to the scene categorization algorithm, except that the spatial properties of the scene are incorporated. For example, regions of a scene that show as bare earth might be either dirt roads or farms. Each of those has distinct spatial properties. A region of disturbed earth may appear to have the spectral shape of bare earth, but will have different spatial characteristics and can therefore be identified as anomalous. This is useful because it can identify potential regions of interest in a complex scene that might be missed by visual inspection of the scene categories map (from the second method).

It should be noted that application of these detection algorithms to specific sites is presently an iterative process. Detection thresholds need to be adjusted for the properties of the location being imaged to achieve consistent and robust results, and comparison of the results from the different algorithms was used to provide insight into interpretation of the results and refinement of the algorithms for this application. That new understanding is the basis for future work to develop more sensitive detection algorithms.

4.0 Analysis Process

For each UNE, published estimates of the location of the UNE were evaluated for a best estimate of the location of the UNE. The Landsat data nearest that location was downloaded for several

data sets before and after the event, subject to data availability and the data being sufficiently cloud-free. A 30 km by 30 km region around the best guess location was extracted from the data for primary analysis, intended to be similar to the 1000 km² search region allowed by the CTBT.

5.0 Results

5.1 First Pakistani Test, 28 May 1998

The 28 May 1998 Pakistani nuclear test was chosen as a first case because there is an obvious surface change in the visible imagery, which is probably due to a rock slide proximate to the location of the UNE. In this case, since the visible imagery clearly indicates significant changes in the surface material, spectral analysis should produce similar results.

Landsat data were acquired for two dates before the event (26 April 1998 and 28 May 1998), and 13 dates after the event (13 June 1998 to 7 January 1999). Figure 1 shows Landsat data one month after the event which illustrates that the event location is obvious in visible imagery, although this single image is not sufficient to unambiguously identify those pixels as proximate to the UNE. Comparison to pre-event data, and inspection of detailed visible imagery (right image in Figure 1) show that the bright region in the image was not present pre-event, and that the bright pixels are due to rock slides.

A subset of the data sets were used for change detection – one date before the event (28 May 1998), and 7 dates after the event (13 June 1998 to 19 October 1998). Some of the data were excluded because they included very thin clouds which would increase the apparent pixel variability and reduce the algorithm sensitivity. The most recent data sets were excluded because of concerns that weathering of the exposed material and human activities would also add unwanted changes to the reference data. (Spectral changes due to material weathering were observed using the material categorization algorithm, confirming the concerns here and illustrating the benefit of using and comparing several algorithms.)

The change detection algorithm determined the statistical significance of change for each pixel (in a multivariate sense) between the one data set taken before the UNE and the larger set of taken after the UNE. The spectra of the statistically significant pixels are shown in the left graph in Figure 2 (the black spectra are from the UNE, all other colors are from other regions showing significant change), and the location of those pixels are shown in the right plot in Figure 2. Note that the graph of spectra show two distinct behaviors, which correspond to two physically separate regions on the ground. Inspection of the visible imagery suggests that the changed pixels in the upper left are due to farm activities, possibly tilling the ground or crop growth. The remaining region corresponds to the suspect region in the visible imagery.

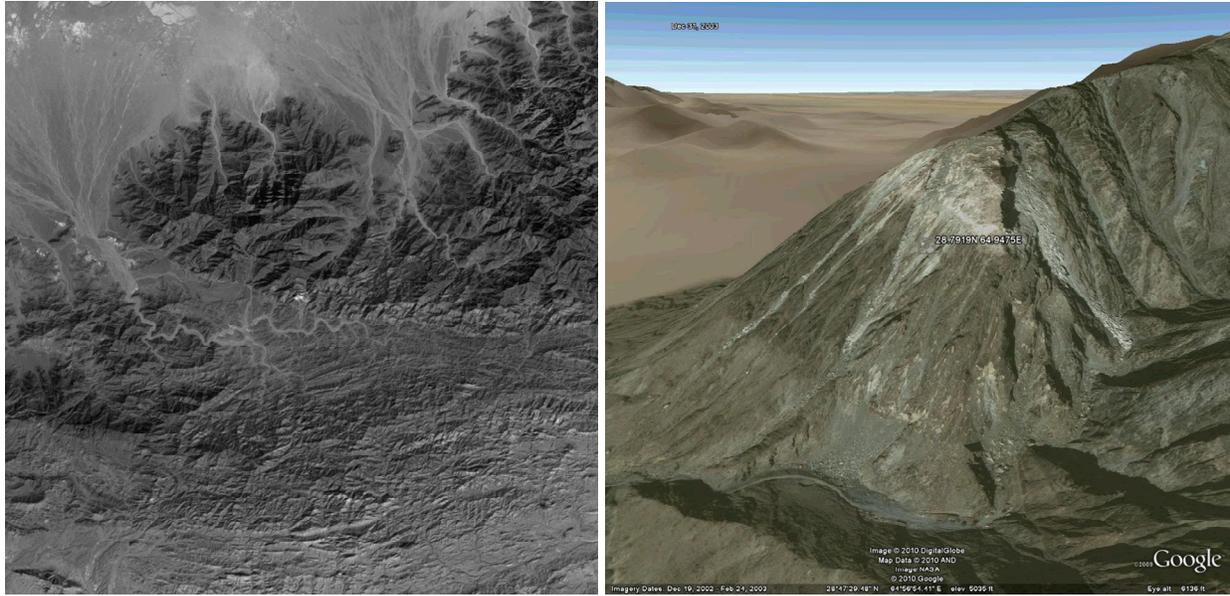


Figure 1. Left: Landsat blue band data for a 30 km x 30 km (1000 pixel X 1000 pixel) region centered on the suspect UNE location (white spot in center). The data is from 29 June 1998, one month after the event. Right: Imagery of that location from 2003, with apparent rock slides down the slopes.

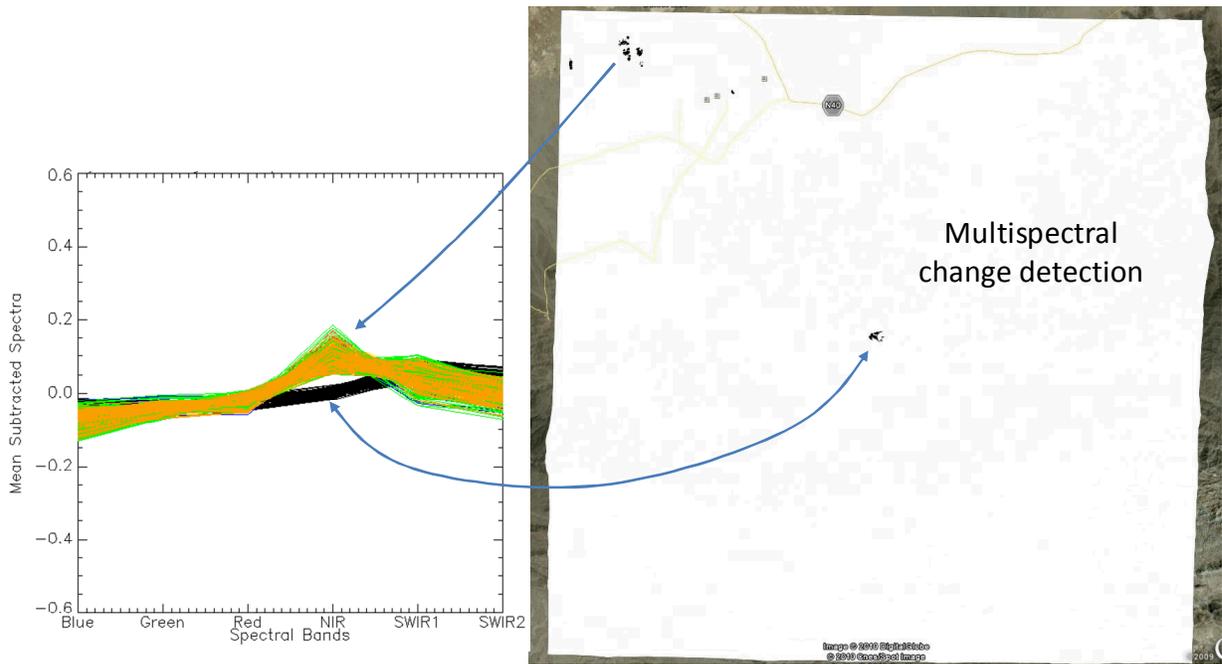


Figure 2. Left: Spectral difference graph shows spectra of pixels identified as having changed significantly after the event. Note that there are two types of spectral behavior observed. Right: The location of pixels with significant change shows them to be in two spatially distinct regions.

The spectral categorization techniques were also applied to this data. Figure 3 shows the results of that analysis. Again, the exposed material from the UNE is readily apparent, and the “false positive” for the same material along the stream is readily distinguished as not relevant. Note that there are no other significant false positives. These results were obtained for each of the 13 post-event images that extend for six months after the event, indicating the spectral persistence of the observable and that no pre-event data is needed to locate the observable. The ability to use only post-event data is directly relevant to an OSI where any MSIR aircraft data will only be collected after the event.

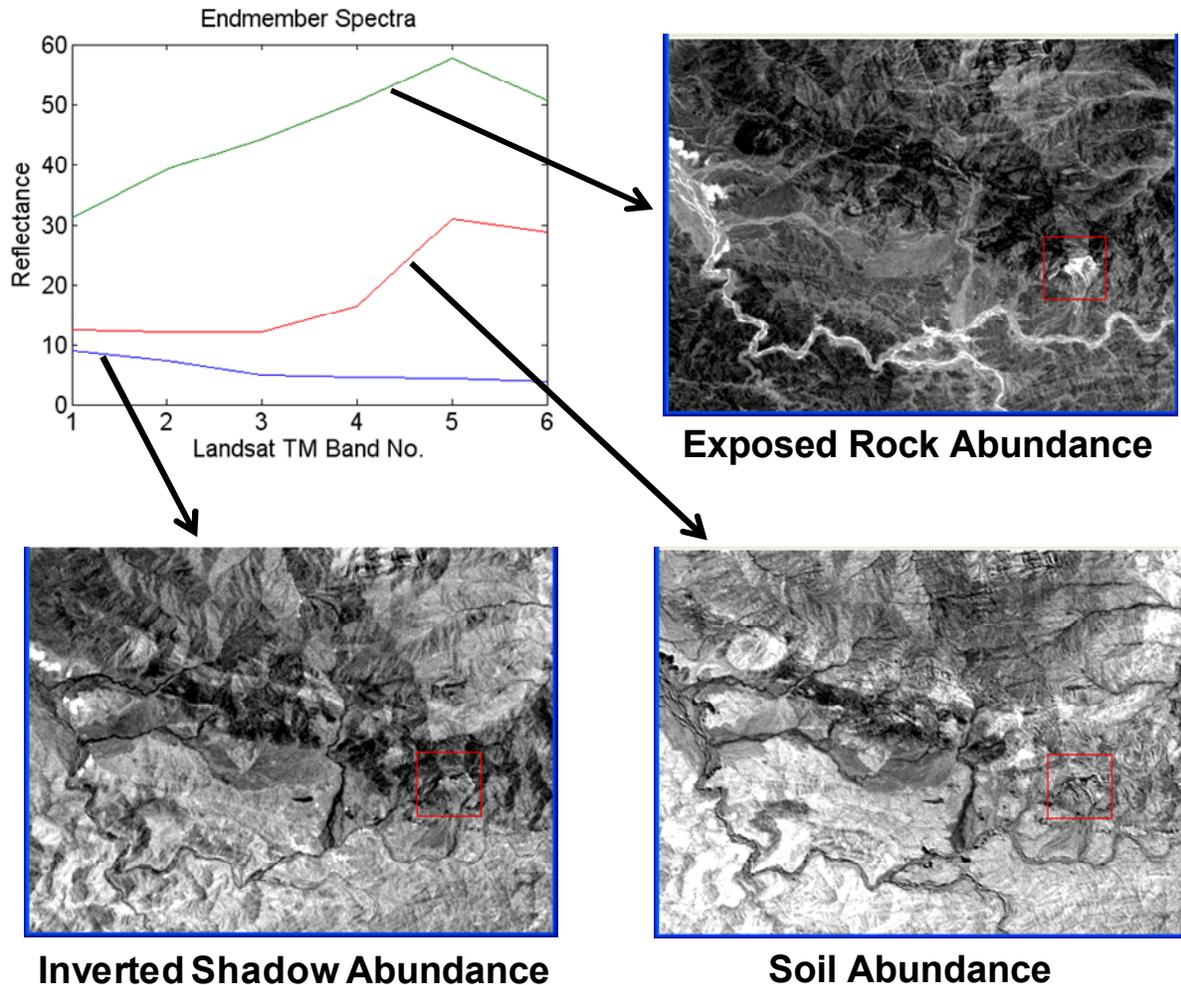


Figure 3. Results of using 6 spectral channels to classify the scene into 3 types of surface material: rock, soil and shadow. Upper left: spectra of the 3 types of surface material. The other images are abundance maps for a portion of the 30 km x 30 km image, showing how similar each pixel on the ground is to one of the three material types. Note that the exposed rock map shows both the region exposed by rock slides from the UNE and material exposed by erosion and roads along the stream.

Here, we also looked at whether the spectrum of the different surface materials changed with time and found that they did for the exposed rock material over the 7-month period after the event. This indicates either that continued rock slides were exposing material with slightly different spectral properties than the original material, or that the exposed material was weathering and its spectral properties were changing over time. In either case, the variation with time of the spectral shape of this region also indicates a region of interest.

While these results are not surprising, given the obvious visible observable, they are important for three reasons. First, it demonstrates that the multi-spectral data and analysis techniques can find observables associated with a UNE. Second, the only “false positives” identified by the techniques were in fact regions one would want to investigate further since they represent disturbed earth (and were readily identified as false positives from the context of the visible imagery). Third, these techniques provide an automated means of sifting through millions of pixels, which is important if higher spatial resolution data is used, and is likely to be essential when rapid and objective data analysis is needed for an OSI.

5.2 Second Pakistani Test, 30 May 1998

The 30 May 1998 Pakistani nuclear test does not have an obvious feature in any visible imagery. Figure 4 shows the 30 km x 30 km regions selected for analysis, along with two estimates of the UNE location, based on seismic analysis with differing assumptions. Two candidate locations for the UNE derived from analysis of the Landsat multispectral data are also shown in the figure for comparison. Thirteen sets of Landsat data were acquired and used in the analyses here, three pre-event (27 April 1990 to 19 May 1998) and 10 post-event (4 June 1998 to 8 May 2000). Again, only some of the Landsat data were used for the change detection algorithm (all three pre-event data sets and only the 4 June 1998 post-event data set).

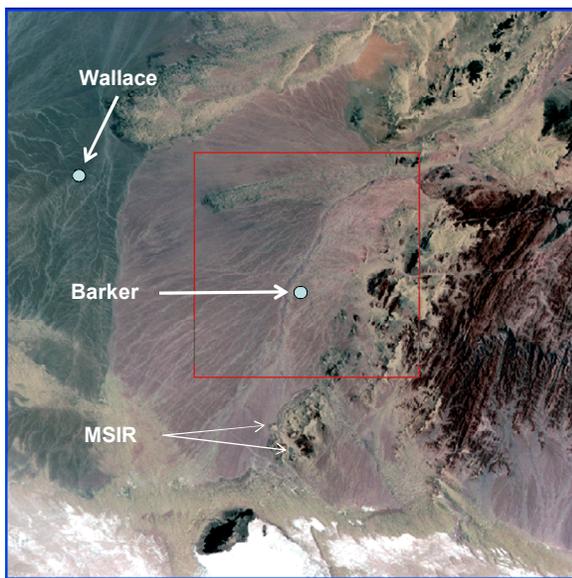


Figure 4. Landsat false color data for a 30 km x 30 km region around the 30 May 1998 Pakistani nuclear test. The test location inferred from seismic data is shown for estimates with two different methods (“Wallace” and “Barker”). Two candidate locations for the test location derived from the MSIR data are shown by the thin arrows. [References: (1) Wallace, Terry C. 1998. "The May 1998 India and Pakistan Nuclear Tests", *Seismological Research Letters*, September 1998. (2) Barker, Brian et al. 1998. "Monitoring Nuclear Tests", *Science*, Vol. 281, 25 September 1998, pp. 1967-68.]

Figure 5 shows the results of the change detection algorithm. Note that the two candidate UNE locations are in a different location than either of the two locations suggested by the seismic data. This is a good demonstration of the limits of accuracy of the ability to geo-locate the UNE seismically for a low yield nuclear test, and of the ability of the multi-spectral data to reduce the search area. The two suspect areas are the regions in the image with the strongest change signal. The other areas in the image showing some significant amount of change are associated with geological features and are likely due to seasonal variations in solar illumination of the scene.

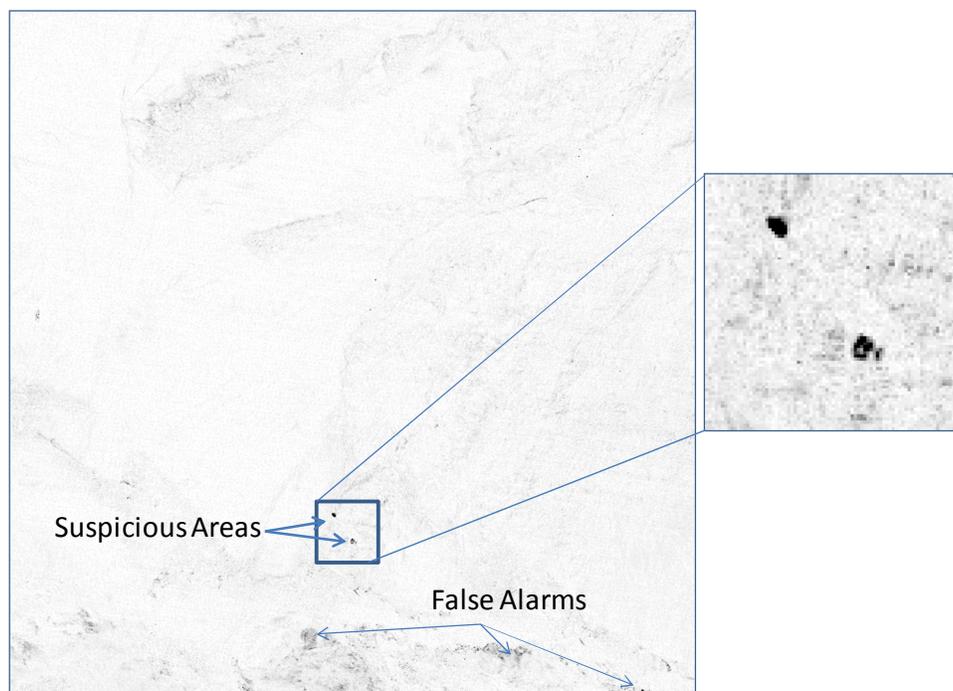


Figure 5. Change detection image for the sets of Landsat data. The degree of statistically significant change is shown in shades of gray, black indicating the most significant change.

Figure 6 shows the results of using the categorization algorithm to distinguish surface materials by their spectral properties. Here the algorithm has been enhanced to include a spatial filter to identify regions where their spatial scale distinguishes them from other parts of the scene with similar spectral properties. This enhances sensitivity and reduces the false positive rate. The results of this analysis are a single high-contrast spectral/spatial region detected for all data on and after 19 May 1998. This is before the test was executed, so the observable is probably due to human activity at that location. Note that this region matches the upper left (North East) suspect region identified from change detection, but that only a single data set of post-event data was used in its detection. The spatial/spectral analysis shows a marginally significant detection for some of the post-event data at the location of the lower right suspect area from the change detection analysis. It is interesting that two very different analysis techniques for the MSIR data give similar results for candidate areas for the nuclear test. It is also not surprising that the

change detection technique appears to be more sensitive, but it is very encouraging that the spatial/spectral technique can find the same regions using only post-event data.

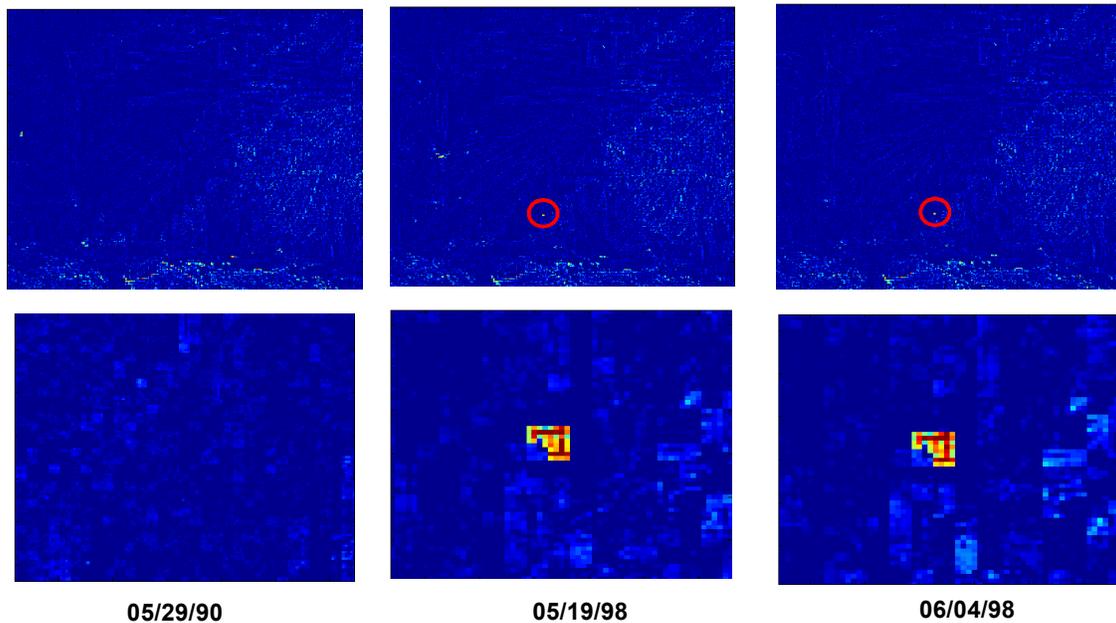


Figure 6. Anomaly detection images (top row) using both spectral and spatial information on single data sets for the dates indicated. The bottom row shows the pixel-by-pixel detail for the circled region in the top row. The results show no anomalies in the data eight years before the test, and show an anomaly for all data sets on and after 19 May 1998. The anomaly here is in the same location as the upper left suspicious area in Figure 5.

5.3 First Indian Test, 11 May 1998

The 11 May 1998 Indian nuclear test does not have an obvious feature in visible imagery. Sixteen sets of Landsat data were acquired, three before the test (9 February 1998 to 14 April 1998) and 13 after the test (16 May 1998 to 21 May 2000). Four of the data sets were used for change detection, three before the test (9 February 1998 to 14 April 1998) and one after the test (16 May 1998). Figure 7 shows the change detection image for the 30 km x 30 km regions selected for analysis based on the best published information for the likely location. Figure 8 shows the region indicated by the change detection analysis superimposed on a current visible image. There is nothing to indicate any surface disturbance in the visible image, but there is about 10 years between the acquisition of the MSIR data and of the visible data. It is interesting to note that the indicated region lies adjacent to an area of obvious human activity.

Spectral image categorization was performed on this data and indicated the same region as a region of interest.

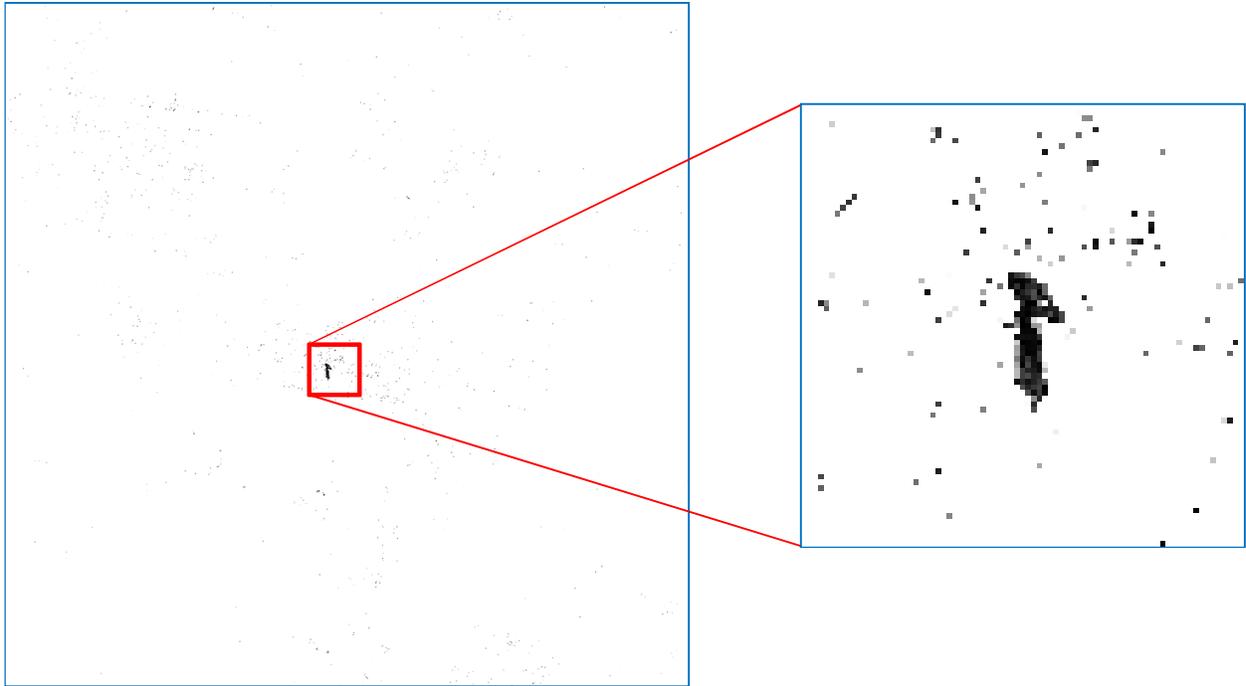


Figure 7. Change detection image for the sets of Landsat data for the 11 May 1998 Indian test. The degree of statistically significant change is shown in shades of gray, black indicating the most significant change. There is only one region of highly significant change.



Figure 8. Visible image showing the location (red outline) of the feature determined from the change detection analysis. There is no obvious visible feature which corresponds to the change region. It is close to an area of human activity.

5.4 Second Indian Test, 13 May 1998

The 13 May 1998 Indian nuclear test was of lower yield than the 11 May test, and public reporting by India indicated it was about 10 km away from the 11 May test. Here we use the same imagery as before, and lower the threshold for change detection and inspect the area about 10 km distant from the first test for candidate locations. Figure 9 shows that there are two candidate regions for inspection that fit the constraints. Since there is no corroborating information for these regions (no visible data, and the image categorization algorithm has not yet been run for these regions), they should be considered as they would for an actual OSI – regions of high interest for early inspection, but not necessarily a definitive location.

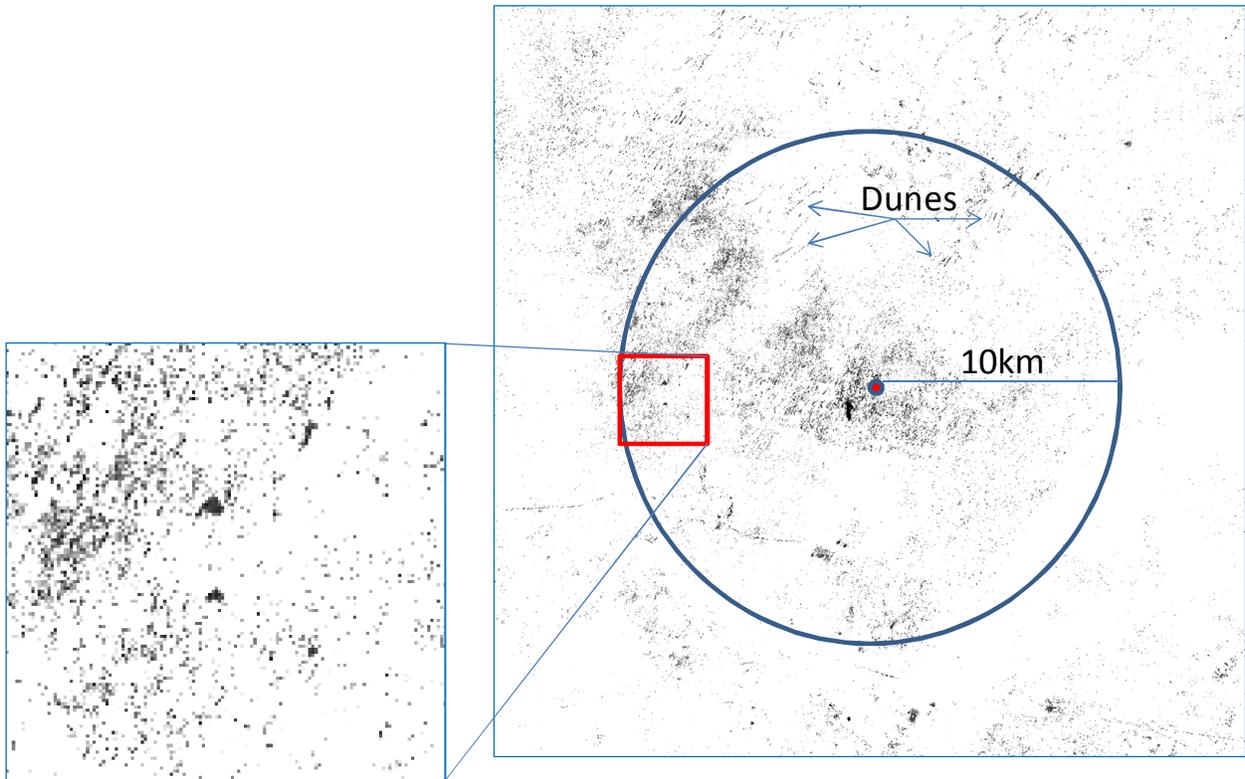


Figure 9. Change detection image for the sets of Landsat data for the 13 May 1998 Indian test. The detection threshold has been turned down to attempt to detect features from the lower yield test. The higher yield 11 May test feature is near the center of the data shown by the red dot. The circle indicates the publicly announced approximate separation between the 11 May test and the 13 May test. The two features in the red box are candidate locations for inspection for the 13 May test since they show significant change, they are within the stated range of the first test, and they do not correspond to false positives from geographical features.

5.5 North Korean Tests, 9 October 2006 and 25 May 2009

5.5.1 Landsat Data for North Korean Tests

The North Korean nuclear tests illustrate some of the limitations of using satellite data. There are frequently clouds in this region, so the Landsat data is often unavailable. For both of these tests, there were only 7 Landsat data sets readily available, but there is a nearly two-year gap with no data around 2008 due to weather and technical problems with the satellite. The data sets are from:

- 22 September 2006 (before 1st test)
- 24 October 2006 (after 1st test, before 2nd test)
- 24 August 2007 (after 1st test, before 2nd test)
- 25 September 2007 (after 1st test, before 2nd test)
- 29 August 2009 (after both tests)
- 29 September 2009 (after both tests)
- 30 September 2009 (after both tests)

Some partially clouded Landsat data may also be available, but adapting the algorithms to account for clouds in some of the data was beyond the scope of this initial effort. It is important to note that the presence of clouds can introduce false positives (detections unrelated to the UNE) for the change detection algorithm, and false negatives (missed detections) for the spectral/spatial algorithm. Also, Landsat 7 data is available for some of this time, but there were problems with the scan mirror on that satellite (the data used here is from Landsat 5), which would have required additional work to attempt to correct for the pointing problems and would likely have resulted in artifacts in the change detection algorithms due to imperfections in pixel co-registration.

Because of the geography and significant plant cover in the region, the North Korean data has larger seasonal variability than the other data analyzed (see Figure 10). This is manifested as varying amounts of snow on the ground, seasonal variations in the vegetation (due to natural variations, logging, and farming), and variations in the shadowing in the imagery due to the mountainous terrain combined with seasonal changes in solar illumination, which impacts pixel signal levels and surface material categorization in the images.

There are two paths that might prove fruitful for demonstrating MSIR utility for the North Korean tests. First, algorithms might be developed that can use partially cloudy data and account for the seasonal affects observed. Second, other satellite data sources might be used, which would mitigate the data availability problem, and might provide inherently more useful data depending on the spatial and spectral resolution. We have started investigating this second path, as described in the next section.

5.5.2 GeoEye-1 North Korean Data

We acquired GeoEye-1 multi-spectral imagery of the North Korean test area with an acquisition date of 12 October 2009, about five months after the second North Korean nuclear test. This is relatively high spatial resolution (approximately 3 m) with a nearly 30 km swath width. (These are about twice the nominal satellite values of 1.6 m and 16 km swath width since this imagery was apparently acquired at a significant slant angle.) GeoEye-1 provides panchromatic imagery and 4-band multispectral imagery (blue, green, red, and near-infrared) with four times the pixel size as the panchromatic imagery. We ran the spatial/spectral algorithm to detect anomalies and identified several facilities of potential interest. While these anomalies are apparent from visual inspection of a color presentation of the imagery (subset shown in Figure 10), the difficulty is that visual inspection results in hundreds of features of comparable apparent significance, often small buildings or clearings in the trees. The spatial/spectral algorithm reduces this number by a factor of 100 or more.

Figure 10 shows Google imagery from February 2005 and GeoEye-1 multispectral imagery from October 2009. The first thing to note is how starkly different the two images are, even though they are both nominally true color images. The difference is that the February 2005 data is in winter with bare ground and bare trees. The October 2009 data still has vegetation present. These large seasonal differences make change detection very difficult with this location. The bright material is probably mining tailings, and is the spatial/spectral observable that caused this portion of the 30 km x 30 km image to be flagged as a region of interest.

It is important to note that this analysis of the GeoEye-1 data used only a single post-event data set to find anomalies, and that the spatial resolution is approaching that expected of airborne MSIR measurements, where sub-meter spatial resolution is readily achievable. Further, the spatial/spectral algorithm highlighted regions of interest automatically, so it was not necessary to evaluate the entire data set, only the regions-of-interest identified.

6.0 Discussion and Path Forward

These early results of the analysis of Landsat Thematic Mapper data show that MSIR observables can be detected from overhead measurements, and that they can be correlated with the likely location of an underground nuclear test. In some cases this is due to disturbed earth, and in other cases it appears human activities generated the observable. It is also very encouraging that two of the early algorithms used here do not require pre-event data, and hence are relevant to the OSI problem where aircraft data will only be acquired post-event. Also relevant is that the regions of interest are few in number and roughly 1 km² in size or smaller, which is a significant reduction from the nominal 1000 km² initial search area.



Figure 10. Comparison of two sets of color imagery of the North Korean test area, before (left) and after (right) the two tests. The left image was taken in February in the middle of winter with the trees and ground bare, and the right in the Fall when the trees and the ground were still vegetated. The white material in the right image appears to be mining tailings. The right image is a false color image generated from the red, green, and blue channels of the 4-channel spectral data from GeoEye-1.

The Landsat data were only likely to find disturbed earth observables, and large-scale human activity observables because the 30 m pixel size is only likely to detect features on that scale or larger. While plant stress might be observed at spatial scales comparable to the disturbed earth observables, it may also require 0.5 m spatial resolution to separate the vegetation from the underlying ground materials. Similarly, many human artifacts would not be resolved with 30 m data. The GeoEye-1 data show that higher spatial resolution is useful for detecting human activities.

Despite these limitations, this work has shown that MSIR can find useful observables with several detection algorithms and a low false positive rate, and that the resulting regions of interest can be used in conjunction with high-resolution visible imagery to provide further discrimination between regions warranting further investigation and those unlikely to be relevant, such as obvious farmlands. The fact that the change detection and the spatial/spectral algorithms identify the same regions using different techniques and different subsets of the data provides further confidence in these results.

The GeoEye-1 data show that higher spatial resolution is useful for detecting human activities. However, this data had only 4 spectral channels, and more spectral channels with higher spatial resolution may provide additional benefit for identifying regions of interest and eliminating false positives. We have started to look at sets of GeoEye-1 and other comparable satellite data (such as that from QuickBird) with the hope that the time series of data can be used to identify those features which are of highest interest, and that information can in turn be used to refine the search parameters so the high-interest regions can be more reliably identified from a single data set.

Finally, it is important to note that an important concern for the maturation of MSIR is to develop the requirements and deployment CONOPS for the MSIR instrumentation that might be flown as part of an OSI. The improvement in ability to identify regions of interest in going from 30 m to 3 m spatial resolution is dramatic, despite the reduction of spectral channels from 6 to 4. This shows that it is likely that aircraft data will have significant utility for reducing the search area in an OSI, and that use of other data with higher spatial resolution and more and different spectral channels might go a long way toward developing the MSIR instrument requirements for an OSI. In the absence of a nuclear test where a post-event aircraft overflight with an MSIR instrument is allowed, a combination of satellite data on UNE's and other events (e.g. mining blasts or earthquakes) along with aircraft data on surrogate events or activities will have to be used to develop the MSIR instrument requirements.

Acknowledgments

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.