Video-rate visible to LWIR hyperspectral image generation and exploitation

Mark Dombrowski

Surface Optics Corporation
11555 Rancho Bernardo Road, San Diego, CA 92127-1441

Paul Willson, Ph. D.

U.S. Army Armament Research and Development Engineering Center
SMCAR-QAH-T, Picatinny Arsenal, New Jersey 07806-5000

ABSTRACT

Hyperspectral imaging is the latest advent in imaging technology, providing the potential to extract information about the objects in a scene that is unavailable to panchromatic imagers. This increased utility, however, comes at the cost of tremendously increased data. The ultimate utility of hyperspectral imagery is in the information that can be gleaned from the spectral dimension, rather than in the hyperspectral imagery itself. To have the broadest range of applications, extraction of this information must occur in real-time. Attempting to produce and exploit complete cubes of hyperspectral imagery at video rates, however, presents unique problems for both the imager and the processor, since data rates are scaled by the number of spectral planes in the cube. MIDIS, the Multi-band Identification and Discrimination Imaging Spectroradiometer, allows both real-time collection and processing of hyperspectral imagery over the range of 0.4 μm to 12 μm.

Presented here are the major design challenges and solutions associated with producing high-speed, high-sensitivity hyperspectral imagers operating in the Vis/NIR, SWIR/MWIR and LWIR, and of the electronics capable of handling data rates up to 160 mega-pixels per second, continuously. Beyond design and performance issues associated with producing and processing hyperspectral imagery at such high speeds, this paper also discusses applications of real-time hyperspectral imaging technology. Example imagery includes such problems as buried mine detection, inspecting surfaces, and countering CCD (camouflage, concealment, and deception).

Keywords:

Hyperspectral-imaging, real-time, visible, infrared, processing, spectral, correlation

Part of the SPIE Conference on Internal Standardization and Calibration Architectures for Chemical Sensors • Boston, Massachusetts • September 1999
SPIE Vol. 3856 • 0277-786X/99/$10.00

1.0 DEMANDS OF HYPERSPECTRAL IMAGING

Real-time imaging over broad bands in the electromagnetic spectrum from the ultraviolet (UV) through the infrared (IR) has been a staple in the areas of remote sensing, surveillance, target detection and tracking, search and homing devices, spectrally tailored coating development, nondestructive inspection, and noninvasive diagnosis. Improvements are being made in these techniques all the time, with increased resolution, higher sensitivity, and greater information throughput being the benefit. The development of high speed digital processing hardware, first at the computer workstation level and later at the integrated circuit level, has spurred the development of more effective optical imaging systems, in part because it has taken up some of the burden formerly handled by analog devices.

Digital image processing paved the way for systems that not only view but analyze the scene in real time. By real time is meant "at video rates", or equivalently, at a rate such that the record of the scene evolves continuously before our eyes, with no gaps or jumps between successive snapshots or frames. Since the human visual system filters input information with a time constant down to approximately 0.1 sec, depending on light level, corresponding conservatively to a noise equivalent bandwidth of 5 Hz, then by the Nyquist criterion essentially maximal information transfer to the brain occurs at 10 Hz and above. In practice, visual perception of a "flicker effect" persists to somewhat higher frequencies, so one might set the threshold frame rate for real-time operation at approximately 20 Hz.

Such a system records the scene in terms of three of the independent variables in which it can be described: x (or azimuth), y (or elevation), and time. It may thus be termed a "3-d" system.

Collecting spectrally resolved information over the broad band viewed adds a fourth dimension to our record of the scene, and provides a basis for a more discriminating analysis of the scene content. Associated with each pixel in the 3-d space defined above is a spectral breakdown of the arriving radiation for that pixel, divided into narrow, normally but not necessarily contiguous bands, several to hundreds in number. In modern parlance such a system is known as *multispectral* (few to several bands) or *hyperspectral* (many bands - which we might define here as a sufficient number of bands to record *all* the *useful* spectral variations to be found in the scene). Systems built on these principles become "4-d" imaging systems, sampling the scene in terms of wavelength as well as the other three variables. As shown in figure 1-1, any hyperspectral imager, save for tomographic projection imagers, use time to generate at least one of the three imaged dimensions. In systems other than MIDIS, however, these cubes are not generated at video frame rates and time is used simply as a parameter to

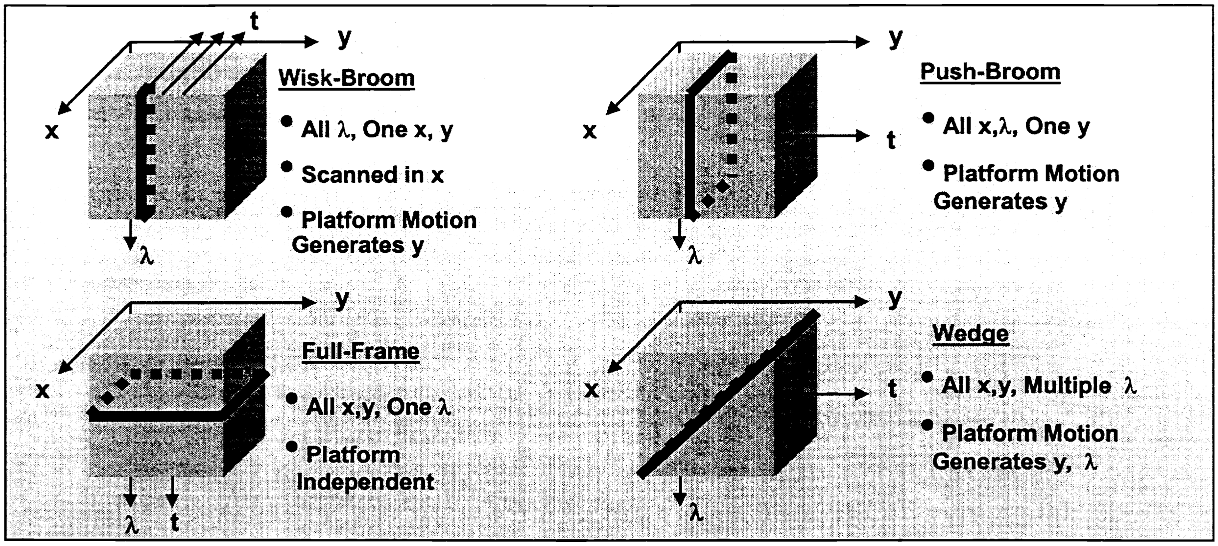

Figure 1-1 Almost all hyperspectral imagers use time as parameter to generate at least one of the three imaged dimensions, x, y, and λ

generate one of the dependent variables, x, y, or λ.

As an outgrowth of the explosion of imaging technology mentioned at the outset, a number of diverse multispectral and hyperspectral imaging systems have been developed and fielded in recent years. The amount of information generated by these 4-d systems becomes so voluminous that their application to many practical tasks depends on the ability to process the data and make decisions on it in real time, that is at rates keeping pace with the real-time acquisition of images. The area of remote threat sensing and reconnaissance includes many such problems. At this point we clarify our definition of real time by asserting that one spatio-temporal frame, complete with all of its spectral points, should be collected at 20 Hz or more. For example, a system operating at 20 spatio-temporal frames per second, collecting 50 spectral bands for each of these, for a total of 1000 *images* per second, may be said to be real-time.

Since 1990, SOC has been engaged in the development of a real-time hyperspectral imaging system, known as MIDIS (multi-band identification and discrimination imaging spectroradiometer). With a combination of Government and in-house resources, a highly capable system has emerged and has been demonstrated in various application areas. A unique aspect of MIDIS is its ability to process the images *spectrally* in *real time*, from initial gain and offset compensation and spectroradiometric calibration through decision oriented processing including spectral correlation using varied algorithms and sensor emulation by spectral integration of the response of other sensors. The uniqueness stems from SOC not only being the first to develop a system collecting and processing spectral images at a 1,000 Hz rate, using custom developed processing boards, but also from the fact that the system, specifically its real-time spectral acquisition and processing capability, is covered by four U.S. patents. Figures 1-2 and 1-3 show the power of spectral filtering techniques to reveal otherwise hidden objects.

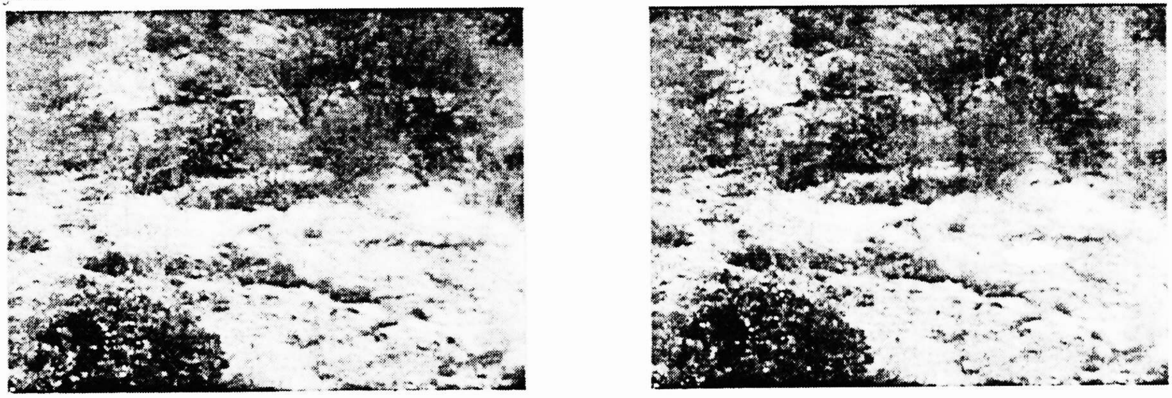

Figure 1-2. Camouflaged Model Tank (left) viewed by VIS-NIR Instrument, Revealed by Real-time Correlation (right).

Figure 1-3. Spectral Correlation works equally well on well-resolved, poorly-resolve, or even sub-resolved targets.

As illustrated above, hyperspectral image data is only useful for the information that it contains. Extraction of the information, however, given the tremendous volume of data that must be processed is indeed a challenge. Information extraction has typically been a job that has taken weeks or months of processing on a workstation, and this for only a few image cubes. Such extended processing times are adequate to analyze data taken from static scenes. But to have the greatest range of applications, e.g., for this technology to find applications on a production line, in the operating room, or on a military vehicle, data collection and exploitation must take place in real-time.

This difference in approach is much akin to the difference between snapshot pictures and motion pictures. Whether collected by exposing film or using CCDs, the image generation technique is the same for snapshot and motion pictures. Only the speed at which the imagery is collected and displayed is different. The information content of motion pictures is significantly greater than that of snapshots, and along with this greater information content comes significantly greater numbers of applications. So it is with hyperspectral imaging. Most systems today are indeed snapshot hyperspectral imagers, capturing a view of a scene at a given instant in time, with no ability to continuously resolve temporal scene changes. Only MIDIS can truly be termed a hyperspectral video camera, capable of continuous real-time collection, exploitation, and presentation of hyperspectral imagery.

The demands on such a system are tremendous. Take for example the aforementioned cube of data comprising 50 images generated 20 times per second, or 1,000 images per second. If each image comprises 256 x 256 spatial samples of the scene, then the pixel generation rate is over 65 million pixels per second. If each pixel, after calibration, is resolved to 16 bits, then the total data rate is over 130 megabytes per second. This sustained data rate is beyond the ability of simple processors to handle, especially when considering that tens to hundreds of operations must be performed on each pixel in the cube to effectively extract the information. Even highly paralleled multi-processor digital signal processors require an excessive number of processors, translating to large size and high power consumption, to accommodate such a rate.

This is the problem that Surface Optics Corporation has sought to solve over its last seven years of involvement in the field of hyperspectral imaging. The resultant MIDIS system is unique not only in its ability to generate high-speed spectral imagery, but also in its ability to extract information from the spectral domain in real-time. In fact, the throughput rate of the processor, the MIDIS Millenium, is a full 320 megabytes per second, nearly a third of a gigabyte per second. With the memory and calibration section of the system occupying a single full-sized PC card, the pre-processor a second card, and the main processor a third, this massive computational power is extremely compact and low-power. The following sections describe this system, its performance, and its applications in further detail.

2.0 MIDIS FUNCTIONAL DESCRIPTION

Figure 2-1 presents a block functional diagram of a MIDIS system with complete coverage from 0.4 μm to 12 μm, delineating both the currently available portions of the system, and those under development. VIS/NIR, SWIR/MIWR, and LWIR spectral imagery is generated at up to 1000 frames per second (per head) by a prism-based spectral imager. Raw imagery data is fed to a memory and calibration board, which uses two ping-pong image memory banks to allow continuous real-time collection and processing of hyperspectral imagery. The memory and calibration board also contains a calibrator circuit with associated calibration memories to translate the raw imagery to true spectroradiometric form. Calibrated imagery is then sent from two or more heads to the input multiplexer and re-mapping board, which selects which, and in which order, bands are sent to the pre-processor. Using a bilinear, random input address re-mapper, this board also corrects for any distortions or mis-registration from head-to-head. Re-mapped calibrated imagery is then sent to the pre-processor, part of the same PC board, which corrects the calibrated apparent scene radiances for atmospheric effects, including absorption, scattering, and path radiance, to produce inherent scene radiances.

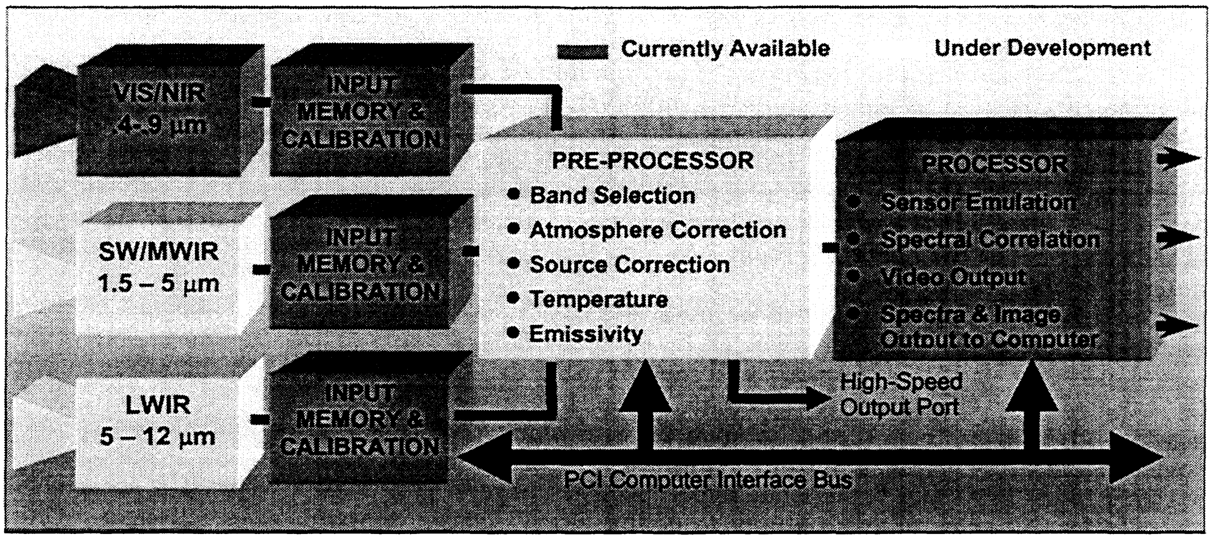

Figure 2-1. The MIDIS system. Hyperspectral imagery is generated by up to three prism-based hyperspectral imagers. A high-speed processor accepts, calibrates, corrects, and exploits the 320 megabyte per second stream of data.

Using JPL, MIT Lincoln Lab, and SOC algorithms, this board also translates inherent scene radiances to scene reflectance or temperature and emittance (depending on the band). Corrected spectral imagery then flows to the main processor, which comprises three independent sensor emulation channels and a spectral correlation channel. Final imagery is output as standard video data, both to the host computer interface and to an external monitor, including a viewer in the imaging head. At any point along the processing chain, data can be sent to the host computer for storage or alternate processing. The high-speed output port from the pre-processor allows data at any level of calibration/correction to be sent to a high-speed mass storage device, such as a RAID device. With such flexibility in data exploitation, MIDIS provides functionality both as a scientific instrument, and as an instrument suited for duty in a tank, an aircraft, in the hospital, or in the factory

MIDIS hyperspectral imagers use a spatially self-scanned prism-based spectral imager. These imagers are capable of collecting spectral imagery at 256 x 256 spatial resolution with variable spectral resolution of 32, 64, and 128 bands. The visible band head also allows collection of up to 256 bands. By using a spatially self-scanned system, the imager is platform independent, capable of use in an aircraft, on a ground vehicle, on an assembly line, or anywhere else. The imager need not move relative to the scene to build up the second spatial dimension as is the case with standard pushbroom imagers. This technique produces a "forward looking" hyperspectral video camera, as opposed to a "down-looking" hyperspectral snapshot camera, analogous to the difference between a FLIR and down-looking IR imager. Imagery is converted to 12 bits resolution, ensuring wide intra-scene dynamic range. Gain of the sensor is variable, both via integration time control and amplifier gain, to allow the system to operate over a broad dynamic range of scene radiances.

Key to the MIDIS system is its real-time spectral image processor. Hosted in a standard PC platform, this processor comprises a three-board set as described above. Available as a stand-alone unit, capable of accepting hyperspectral imagery from any hyperspectral imager – given that the output data is formatted to the Fibre-channel serial input format of the processor – at input rates up to 160 million pixels-per-second, this hyperspectral image processor provides video-rate data exploitation. Data from previous acquisitions can also be streamed into the processor from mass storage devices such as RAID systems, or uploaded from disk through the computer's PCI bus. SOC uses a Dell Precision 410 workstation to house the processor. Figures 2-2 and 2-3 show the Frame Interface and Calibration (FIAC) board and the High-speed Processor and Video Output (HPVO) board, respectively, from the MIDIS hyperspectral image processor.

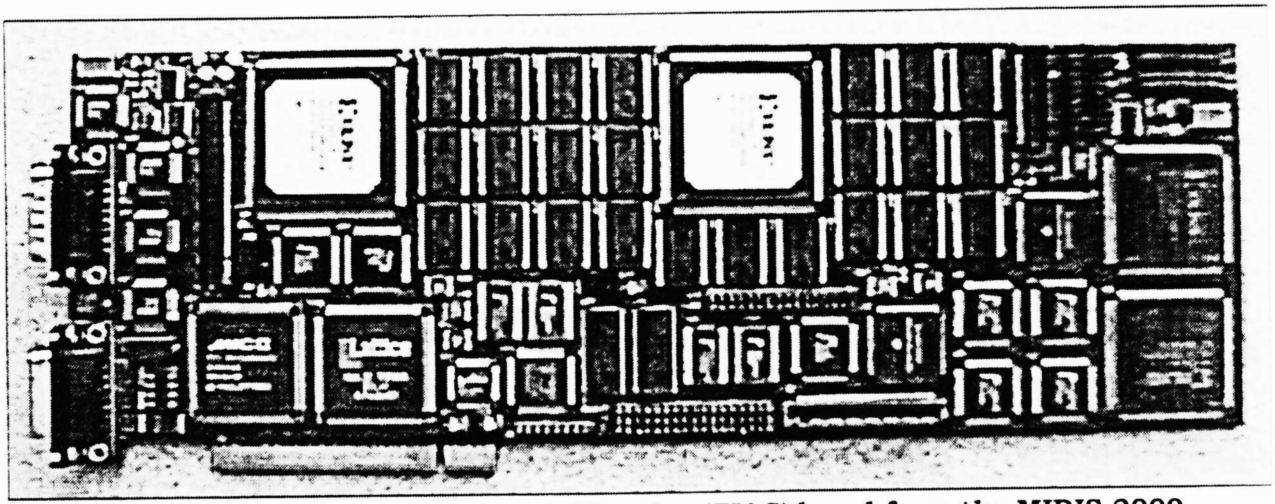

Figure 2-2 The Frame Interface and Calibration (FIAC) board from the MIDIS-2000 hyperspectral processor collects and calibrates imagery at 160 mega-pixels per second

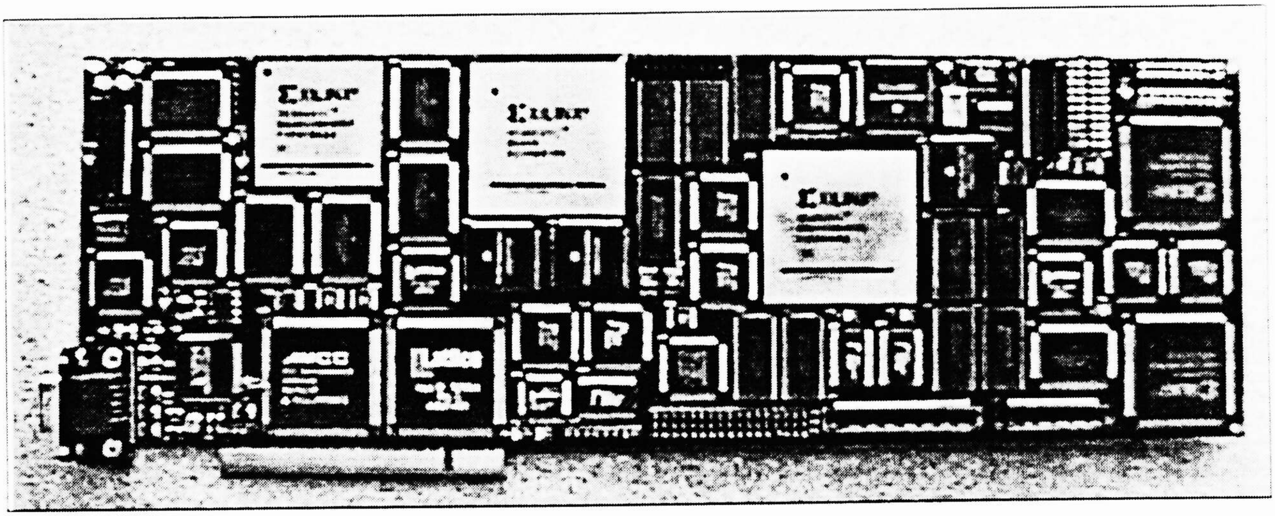

Figure 2-3 The High-speed Processing and Video Output (HPVO) board from the MIDIS-2000 processor performs sensor emulation and spectral correlation in real-time

3.0 EXAMPLE REAL-TIME APPLICATIONS

Presented at the outset of this paper (section 1.0) were two examples of military applications that require real-time collection and processing of hyperspectral imagery, and also require forward-looking imaging capability. These two, extraction of a threat (tank) from background, require that the threats be found now, not two weeks from now. A tank commander who must wait two weeks, or even just hours or minutes to find out if an enemy tank is targeting him will find out much quicker in a much less pleasant way. He must have the information immediately, especially if he is to scan a scene to identify the threats. Similarly, an advancing convoy or patrol must know now where the mines are that the must destroy or avoid.

As tremendous as the military potential for this technology is, the industrial, medical, and commercial possibilities are even larger. For instance, the ability to identify surface contaminants throughout an image would be a boon to a wide variety of applications, including coatings and semiconductor production, aircraft/vehicle painting, anti-terrorism (in the form of identifying explosive compounds on vehicles or packages), etc. Since production or processing capability translates directly to profitability, any system designed for such applications must produce immediate indication of surface cleanliness. Figure 3-1 shows an example of composite coupons viewed by a broad-band IR imager, and the results of filtering spectral imagery to identify areas of mold-release contamination where subsequent coatings would not adhere.

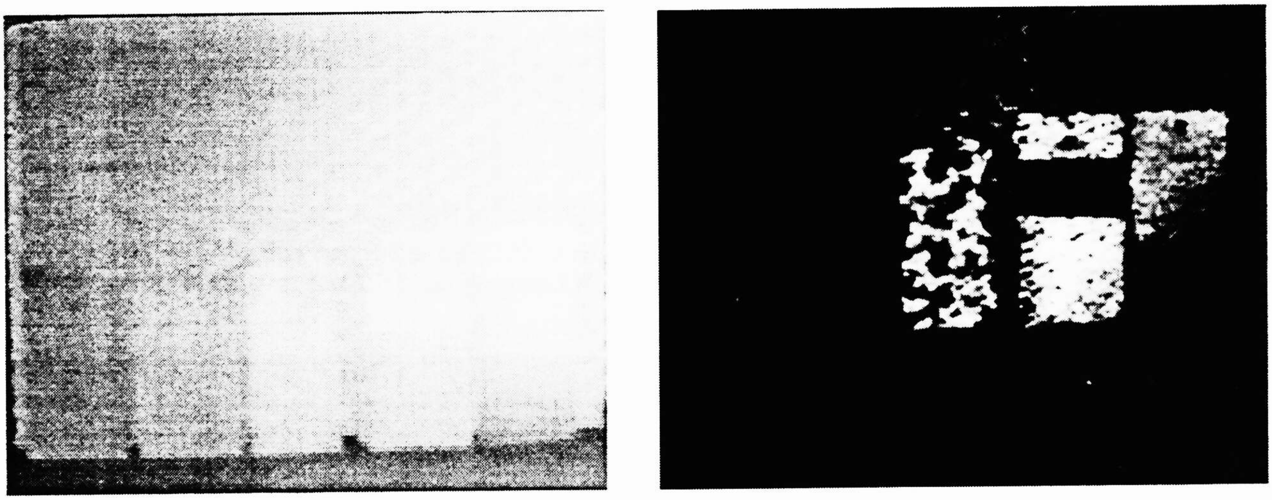

Figure 3-1. Broadband IR imagery (left) fails to show the mold-release contamination on composite coupons. Real-time spectral correlation on hyperspectral imagery (right) reveals the contamination.

Medical applications of real-time hyperspectral imagery are also tremendous, from aiding the pathologist in identifying cancerous cells on slides, to *in vivo* identification of such cells and tumors, to identifying necrotic tissues. Any medical need to differentiate between tissues, whether within the body or without, would benefit from this technology.

Another area of active research is detection of buried objects – buried mines, specifically – through exploitation of hyperspectral imagery. Burying a mine results in a discernable spectral signature difference relative to the undisturbed background. In an environment without foliage, this change occurs due to smaller particles, normally washed or blown away, being brought to the surface. Changes in soil particle packing density also change surface emissivity. Surface Optics obtained soil samples from Murac dry lake in southern California, some carefully taken to avoid disturbing the surface, and others of the same soil but with surface disruption as would result from burying an object. These samples were then measured in our hemispherical directional reflectometer. Figure 3-2 illustrates the change in soil emissivity from an undisturbed surface to a disturbed surface. Note significant difference in spectral characteristics between the undisturbed and disturbed soil, both in the 2.5-5 μm and especially in the 6-12 μm range – ranges covered by MIDIS's SWIR/MWIR head and LWIR head respectively. Coupled with the difference in heat capacity and thermal conductance of the mine relative to the surrounding soil, this emissivity difference gives rise to a spectral signature uniquely identifying the location of buried objects. Note, however, that time will tend to erase the surface emissivity differences, although the period over which the surface will return to an undisturbed appearance is unknown.

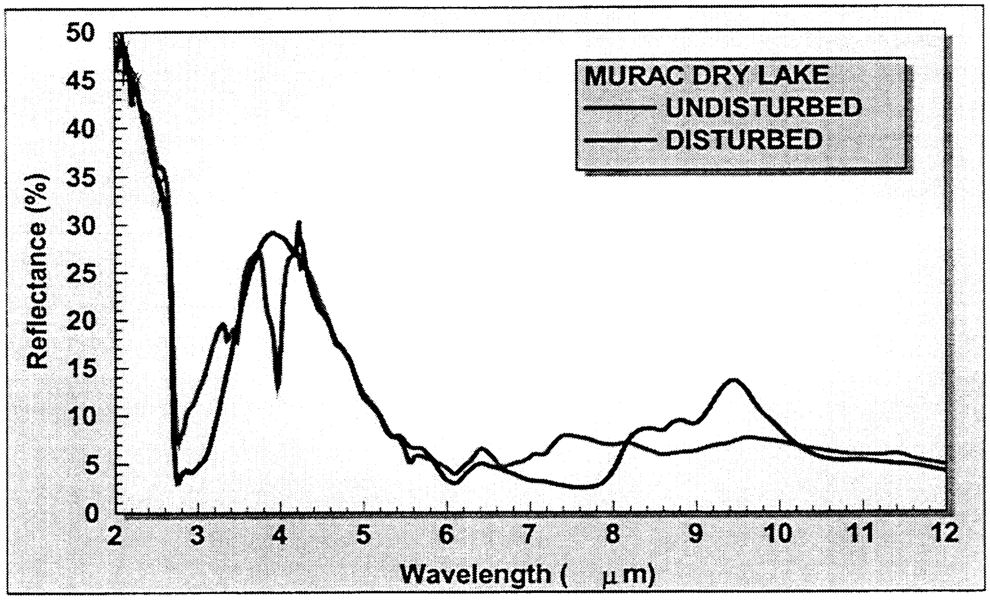

Figure 3-2. Burying a mine brings small particles to the surface, changing surface emissivity and revealing the mine location

A synthetic scene was generated using this emissivity data. Using the undisturbed soil as the main constituent of the scene, mines were embedded in a small array pattern. Vegetation was also dispersed throughout the scene. Soil temperatures ranged from 24°C to 28°C, with plant temperatures running about 2°C cooler.

A cube of data representative of this scene viewed by the MIDIS-2000L LWIR hyperspectral video camera was generated from the scene temperature data and surface emissivity data. For purposes of demonstrating the robustness of the spectral correlation algorithms embodied by MIDIS-2000, SNR was set at only 50:1 across the spectrum, significantly lower than the predicted value shown in section approaching 1,000:1. After generating the cube of data, a broad-band 5-12 μm response curve was applied, representing use of one of the three sensor emulation channels to be included in the system. Figure 3-3 shows the resulting broad-band image.

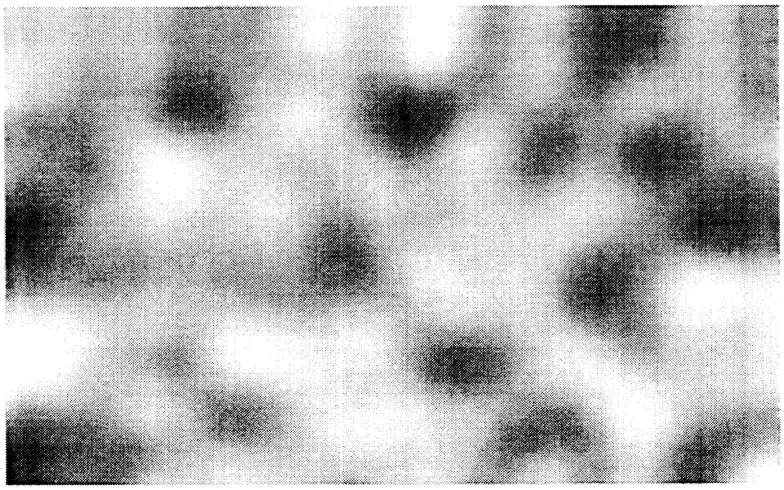

Figure 3-3. Broad-Band image of synthetic scene developed to investigate buried-mine detection

Embedded in the above scene are 12 mines in two regular arrays. Because of the clutter in the scene, and lack of broad-band contrast of disturbed soil to undisturbed, the mines are not visible. Figure 3-4 shows the spectral radiance of three different pixels in the scene, one taken from undisturbed soil, one from vegetation, and one from disturbed soil. Atmospheric transmission loss is not shown.

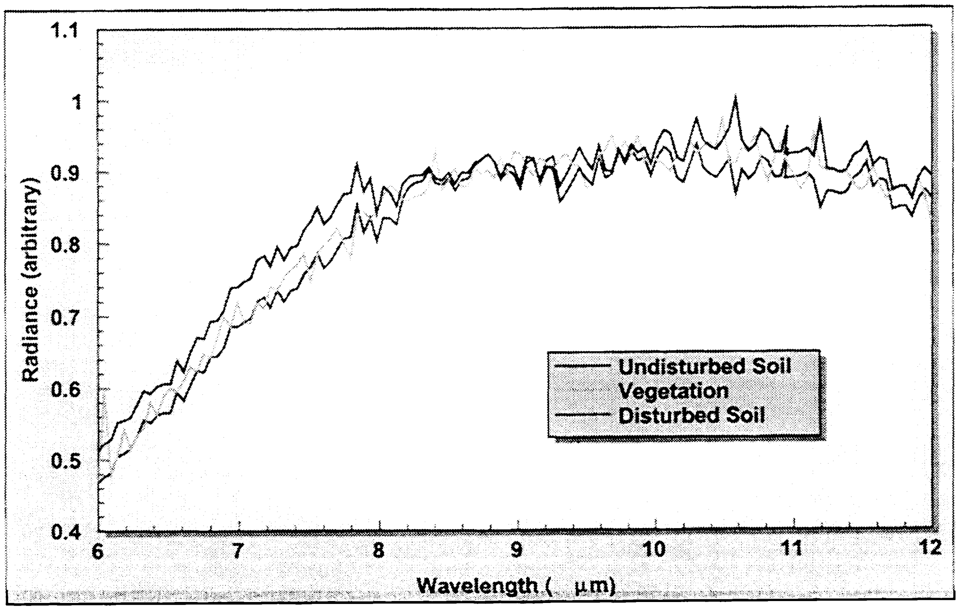

Figure 3-4. Comparison of spectral radiances from the three classes of materials in the synthetic scene. Note the intentionally poor SNR

Note that the spectra show subtle differences representative of their differing emissivities. The cube of synthetic imagery was then processed using one of the spectral correlation algorithms implemented by MIDIS-2000 to identify the areas of disturbed soil. Results of this spectral correlation are shown in figure 3-5. Spectral correlation transformed the broad-band image with no discernable mines into a clearly resolved image of mine locations. This modeling study clearly shows the capacity of the LWIR MIDIS system to resolve objects based on differing material characteristics, even when the differences are very subtle and embedded in relatively large noise.

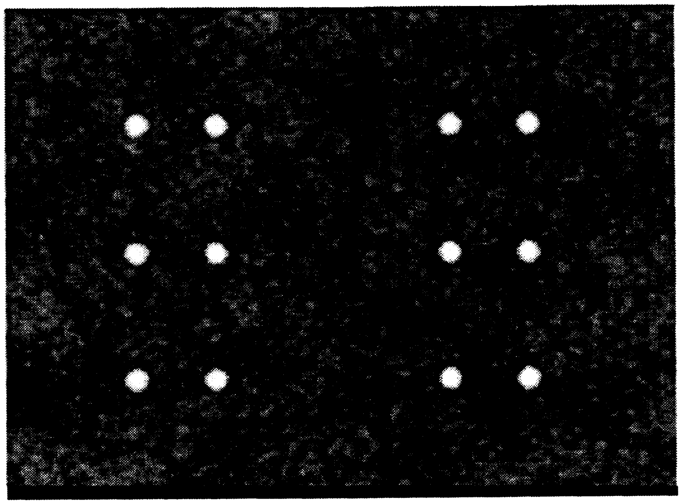

Figure 3-5. Spectral Correlation clearly reveals the buried mine locations

4.0 CONCLUSION

For hyperspectral imaging to realize its potential to provide information in an image based not merely on contrast differences, but more importantly on material differences revealed through each pixel's spectrum, ability to process in real-time is absolutely necessary. A military threat sensor is of no use if it reveals the threat days, hours, or even minutes after viewing a scene. Every second delay in revealing a threat gives that threat another second to act. In medical applications, used either in a doctor's office or pathologist's lab to identify tissues, e.g., malignant versus benign, processing must keep up with the doctor's ability to search for problems. In an operating room, used to guide the surgeon's hand during excision of malignant, necrotic, or other tissue, the processed hyperspectral imagery must present seamless, real-time motion. On the manufacturing line, where speed of production equates to profitability, the HS imager must be able to identify improperly manufactured materials – missing coatings, improperly mixed chemicals, improperly formed polymers, etc. – at the same rate that they are produced. The list of potential real-time HS imaging applications is tremendous.

Only by fulfilling this need for real-time information will hyperspectral imaging move from being a specialty imaging technology used mainly as a research tool to being a ubiquitous imaging technology used to solve numerous and diverse problems. Surface Optics Corporation has devoted considerable effort and resources to advancing hyperspectral imaging to hyperspectral video. This effort has produced a new tool available to aid in timely processing of hyperspectral imagery – the MIDIS-2000P hyperspectral image processor. The MIDIS-2000V, -2000M, and –2000L, Vis/NIR, SWIR/MWIR, and LWIR real-time hyperspectral imaging heads generate spectral imagery fast enough to exploit the processors speed, creating a true hyperspectral video system. Considerable work remains to be done to bring this technology to its ultimate fruition, both in continued size and cost reduction. Surface Optics is currently actively addressing these items.

5.0 ACKNOWLEDGEMENTS

Surface Optics Corporation thanks Dr. Paul Willson at US Army ARDEC for his continued support in furthering development of the technology originally devised and patented at SOC. Without his support, real-time HS imaging would not be advanced to the state that it is today.