

Vehicle tracking with multi-temporal hyperspectral imagery

John Kerekes^{*}, Michael Muldowney, Kristin Strackerjan, Lon Smith, Brian Leahy
Digital Imaging and Remote Sensing Laboratory
Chester F. Carlson Center for Imaging Science
Rochester Institute of Technology
54 Lomb Memorial Drive
Rochester, New York, USA 14623

ABSTRACT

Hyperspectral imagery has the capability of capturing spectral features of interest that can be used to differentiate among similar materials. While hyperspectral imaging has been demonstrated to provide data that enable classification of relatively broad categories, there remain open questions as to how fine of discrimination is possible. An application of this fine discrimination question is the potential that spectral features exist in the surface reflectance of ordinary civilian vehicles that would enable tracking of a particular vehicle across repeated hyperspectral images in a cluttered urban area.

To begin to explore this question a vehicle tracking experiment was conducted in the summer of 2005 on the Rochester Institute of Technology (RIT) campus in Rochester, New York. Several volunteer vehicles were moved around campus at specific times coordinated with over flights of RIT's airborne Modular Imaging Spectrometer Instrument (MISI). MISI collected sequential images of the campus in 70 spectral channels from 0.4 to 1.0 microns with a ground resolution of approximately 2.5 meters. Ground truth spectra and photographs were collected for the vehicles.

These data are being analyzed to determine the ability to uniquely associate a vehicle in one image with its location in a subsequent image. Initial results have demonstrated that the spectral measurement of a specific vehicle can be used to find the same vehicle in a subsequent image, although this is not always possible and is very dependent upon the specifics of the situation. Additionally, efforts are presented that explore predicted performance for variations in scene and sensor parameters through an analytical performance prediction model.

Keywords: Hyperspectral vehicle detection, target tracking

1. INTRODUCTION

Hyperspectral imagery has found application in a variety of fields due to its ability to capture the differing spectral characteristics of materials. Classes of vegetation, minerals, and even man-made construction materials can be distinguished through measurement of their varying spectral reflectance curves. The spectral features leading to this discrimination can come from the materials absorbing radiation at certain narrow wavelengths or reflecting light with a gentle sloping but distinctive change over wavelength due to scattering in the top surface of the object.

In all practical spectral imaging systems the light collected in a given pixel comes from a spatially distributed and heterogeneous area of the surface (ranging from less than a meter to ten's of meters) that encompass many variations in the reflectance of a given material, or even multiple materials in varying compositions. As a result, while many pixels from a given surface class or material may be very similar, they are most likely to be different or even unique in their elemental composition. For example, pixels over a grass field will sample different combinations of healthy grass blades, dying grass, dead grass, decomposing organic matter, weeds, and soil. Variations in the angular geometry of the grass blades, etc., will also lead to varying reflectance. The variations may be small, but they most surely will be present. Variability is similarly present in man-made materials, although usually to a smaller degree, and can come from variation in the production process, uneven weathering as well as contamination from natural elements, etc.

* Contact information: kerekes@cis.rit.edu.

The work described in this paper is part of an ongoing project to assess the feasibility for particular objects of interest to be located and tracked in sequential frames of hyperspectral imagery through the use of their potentially unique spectral reflectance characteristics. In this present work, we are investigating the idea of locating a vehicle in a given hyperspectral image, extracting its spectral characteristics, and then using that information to find the same vehicle in a subsequent image. This concept would require that there be sufficient distinguishing characteristics in the vehicle spectral reflectance so it is not confused with other similar vehicles, and that those characteristics remain present during the interval between the hyperspectral image collections.

This paper presents the results of a recently conducted experiment to explore the feasibility of vehicle tracking with a hyperspectral imager. The following sections describe the experimental data collection activities and preliminary results of the analyses. Also presented are the results of a model-based analysis exploring the parameter sensitivities and tradeoffs in the detection of vehicles in the imagery.

2. EXPERIMENTAL DATA COLLECTION

On 24 June 2005, an experiment was conducted on the Rochester Institute of Technology (RIT) campus involving the use of volunteer students, staff and their vehicles. At pre-arranged times the volunteers positioned their cars at various locations around the campus. Coordinated with the movement between locations, multiple passes over the campus were made by an airborne hyperspectral imager. Spectral reflectance ground truth measurements were made of various vehicles involved in the experiment. The following describes these activities.

2.1. Movement of vehicles

Six volunteers and their vehicles were provided instruction sheets and asked to park their car in a given designated lot on campus and to record their precise location using diagrams provided. After the first pass of the aircraft over their area, they were then instructed to move to a second designated lot and again record their precise location and wait for the second over flight. Some (but not all) locations were documented with digital photographs. Also, some volunteers were instructed to park in open areas while some were not given specific constraints. Most volunteers were diligent in recording enough information to precisely locate their vehicle, but not all and in some cases it would have been better to have had photographs showing the location and immediate surrounds.

2.2. Airborne data collection

The airborne hyperspectral imagery was collected by the Modular Imaging Spectrometer Instrument (MISI) designed and built by students and staff in RIT's Digital Imaging and Remote Sensing (DIRS) Laboratory. MISI was flown at an approximate altitude above ground level (AGL) of 2500 feet in a Piper Aztec making two passes each over the eastern and western sides of the campus. Section 3 provides further details on MISI.

2.3. Spectral ground truth data collection

An Analytical Spectral Devices, Inc., FieldSpec Pro was used to collect spectral reflectance ground truth for some of the vehicles involved in the experiment and for two sets of "calibration" panels used to convert the MISI data to surface reflectance. Not all of the vehicles were measured due to time constraints and logistical issues. Also, some were measured not on the day of the experiment, but later in the summer as part of a separate research effort on environmental effects.

Two different types of "calibration" panels were deployed during the experiment. They are visible near the top of the left side images shown later in Figure 4. The large black and white squares are areas of the outer parking lot (not used in the summer) painted with black and white paint. These squares were approximately 100' x 100'. Directly below these panels are two smaller black and light gray panels that were 30' x 30' in size and made from canvas fabric. The fabric panels were deployed just on the day of the experiment and held down by small rocks on the perimeter, while the painted panels remained in use for subsequent repeat collections throughout the summer and early fall¹.

3. MISI DESCRIPTION

With contributions from staff and students the airborne imaging spectrometer known as MISI has been developed to support of a variety of remote sensing research projects². MISI's development spans nearly fifteen years of evolution.

MISI is a line scanner with a rotating mirror coupled to a Cassegrain telescope with separate focal planes covering the visible through long wave infrared. It contains two imaging spectrometers collecting 70 channels across the visible through near infrared and several broadband detectors spanning the shortwave through longwave infrared. Table 1 provides specifications of the instrument.

Table 1. MISI specifications.

Spectral Band	Center λ	# Channels	$\Delta\lambda$	GIFOV @ 2000' AGL
VIS	0.41 to 0.75 μm	35	0.012 μm	6'
NIR	0.74 to 1.02 μm	35	0.010 μm	6'
SWIR	1.26 μm	1	0.11 μm	4'
SWIR	1.65 μm	1	0.38 μm	4'
SWIR	2.03 μm	1	0.65 μm	4'
MWIR	3.65 μm	1	0.9 μm	4'
LWIR	9 μm	1	2 μm	4'
LWIR	11 μm	1	2 μm	4'
LWIR	11.5 μm	1	2.1 μm	4'
LWIR	11 μm	1	6 μm	2'

The instrument operates in a line scanning mode with a rotating mirror collecting incident light and reflecting it onto three separate focal planes in the cross-track direction. A $\pm 45^\circ$ field-of-view allows collection of 2 km wide swaths from 1 km AGL. During each line scan, the detectors also view visible and thermal calibration sources. Figure 1 shows the instrument undergoing testing in the laboratory



Figure 1. MISI in the laboratory.

4. DATA COLLECTED

4.1. Spectral ground truth

Spectral reflectance measurements were collected with an Analytical Spectral Devices, Inc., FieldSpec Pro for a number of vehicles and the calibration panels. Figure 2 shows the set-up that was used to collect reflectance spectra of vehicles subsequent to the day of the experiment. Figure 3 shows example spectral reflectance curves for a number of different vehicles. The blue and green vehicles show peaks just below and above 500 nm in wavelength. A white vehicle shows a very high reflectance in the visible but falling steadily at the longer wavelengths. As can be seen, nearly all vehicles exhibited quite high reflectance in the near- and shortwave-infrared.



Figure 2. Spectral ground truth data collection platform




Figure 3. Example spectral reflectance (%) for several vehicles.

Just visible in Figure 2 on the right in the background is the black painted calibration panel. During the day of the experiment measurements of this panel, the white painted one, and the black and light gray fabric panels were made to support the conversion to surface reflectance of the airborne imagery. Further details on this atmospheric compensation process are described in Section 5.1.

4.2. Airborne hyperspectral imagery

In order to achieve high spatial resolution, MISI was flown at the lowest altitude possible considering flight restrictions and the speed of the scanning mirror. The 2500' AGL altitude with a 3 mrad instantaneous field of view resulted in the data having a spatial resolution of approximately 7.5 feet. At this low altitude it took two separate flight lines to collect imagery over the entire campus, even with the $\pm 45^\circ$ field of view of the sensor. Portions of the imagery collected for the various passes are shown in Figure 4.

The images on the left side of Figure 4 correspond to the two passes (one at 12:09 pm local time and the next at 12:24 pm) over the western side of campus, while the images on the right side correspond to the two passes (one at 12:04 pm and the second at 12:28 pm) over the eastern side of campus. The running track (near the top) and the tennis courts (near the bottom) are visible in the images from both flight lines indicating the overlap of the swaths.

Several features of the data are obvious from these images. It is clear that the current roll-correction algorithm does not entirely remove the effects of platform roll in the imagery. Straight roads and building edges have some waviness to them. Geometric effects due to the wide scan angle also remain visible. Additional processing would be necessary to ortho-rectify the images.

Also, while the 12:04, 12:09, and 12:28 pm images are seen to be fairly clear, the image collected at 12:24 pm shows a large cloud shadow over the right side. The shadow covers the area of the calibration panels which means the retrieved reflectances will only be reasonable for the parts of the image covered by a similar level of shadow. This is discussed further in Section 5.1.



Figure 4. MISI imagery collected over the RIT campus. The times indicate the local time of each pass.

5. EMPIRICAL DATA ANALYSIS

The MISI images were analyzed by first converting to surface reflectance through an atmospheric compensation process and then applying matched filter detection operators to locate specific vehicles. The empirical data analyses were conducted with the ENVI™ software available from Research Systems, Inc.

5.1. Conversion to spectral reflectance

Spectral measurements of the large (100' x 100') painted panels were used with the Empirical Line Method (ELM) to convert the raw digital counts of the MISI images to surface spectral reflectance. Regions of Interest (ROI's) were used to select pixels from the imagery over the panels and then a wavelength-dependent linear function was developed between the known ground reflectance of the panels and the MISI data. These functions were then applied to all pixels in the images to convert to surface reflectance.

This procedure was used for the 12:09 and 12:24 images, but since the calibration panels were not in the 12:04 and 12:28 images, other areas were selected. Portions of the running track (near the top) and the artificial turf field (near the bottom) were used for these two images. Since ground measurements of these areas were not readily available, the results of the atmospheric compensation applied to the 12:09 image were used to provide the surface reflectance for the track and turf field.

Figure 5 presents example single-pixel spectra resulting from the ELM atmospheric compensation process applied to the 12:09 pm MISI image. While there is some noticeable noise, particularly in the 950-1030 nm spectral region, the major spectral characteristics of the different objects are clearly visible and maintained in the data.

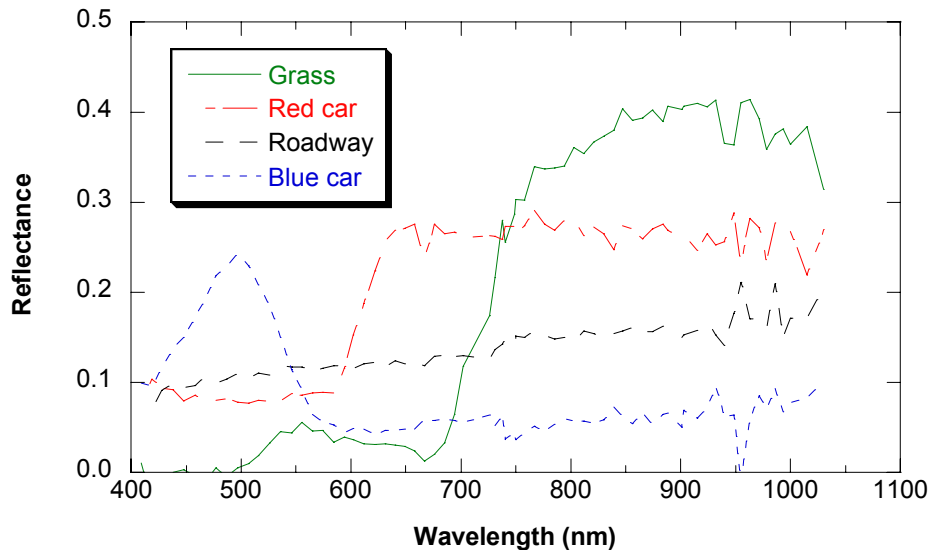


Figure 5. Example reflectance (0-1) spectra resulting from the ELM-compensated 12:09 pm MISI image.

5.2. Finding a vehicle in the same image

The first analysis attempted was to explore the case of using an in-scene measurement (spectrum) of a car to find the car in the same image from which the measurement was extracted. While this may seem like a trivial experiment, the result could reveal fundamental problems with the vehicle tracking concept. If this were not possible, then it would make little sense to continue the effort!

The top image in Figure 6 is a photograph taken from the Chester F. Carlson building on campus showing an array of four blue vehicles spaced well apart in the second row of the parking lot. These four are circled in the zoom of the 12:09 MISI natural color image shown in the lower left of the figure. A pixel from the MISI image was selected from the upper rightmost vehicle (the one circled in the photograph) and used as the target in a matched filter algorithm

applied to the image subset (400 x 500 pixels) as shown in Figure 4. The gray-level matched filter result for the zoom image area is shown in the lower right of Figure 6.




Figure 6. The oval on the zoomed MISI image (12:09 pm) in the lower left encircles the four blue cars shown in the second row of the upper photograph. The lower right image is the matched filter output obtained from using a pixel from the MISI image corresponding to the circled vehicle in the photograph and applying it to the same MISI image. The matched filter output had the highest value of the entire image for the pixel used as the target.

The pixel in the matched filter output with the highest value (=1.00) indeed turned out to be the pixel that was used as the target in forming the matched filter detector. The next highest value in the MF image was 0.54 and it occurred near the bottom of the image where large blue tarps were covering construction materials stored at the edge of a parking lot. The highest output for another vehicle was found to be 0.36, far below the true target's output. This result confirmed that across 200,000 pixels the particular pixel for a vehicle was its own best match and no other vehicle was even close.

5.3. Finding vehicle in second image using spectrum from the first

The second analysis was to explore the case of using an in-scene measurement (spectrum) of a particular car from the first pass image to find the same car in the second pass image. This is the next step in investigating the feasibility of vehicle tracking with hyperspectral imagery. If this could be accomplished with low numbers of false detections, then the feasibility could be established. Figure 7 shows the situation and the result. A pixel was extracted from the 12:09 image corresponding to the isolated blue car shown in the photograph. That pixel was then used as the target spectrum in a matched filter applied to the 12:24 image shown in the lower left of Figure 7. The lower right image in Figure 7 shows the scaled matched filter result.



Figure 7. The circle on the MISI image (12:24 pm) in the lower left encircles the blue car shown in the upper photograph. The lower right image is the scaled matched filter output obtained from using a pixel from the 12:09 MISI image corresponding to the circled vehicle in the photograph and applying it to the 12:24 pm MISI image.

The first observation is there are a significant number of pixels which have a gray level (corresponding to the matched filter output) higher than the circled pixel containing the blue car (its value was 1.26). Upon further quantitative analysis one finds that only six other pixels on cars have matched filter outputs that exceed that of the desired blue car (ranging from 1.28 to 1.41). The rest of the bright pixels correspond to locations on buildings or the blue tarp (maximum in the image of 2.10) covering construction material near the bottom of the image.

Thus, this analysis demonstrates the use of an image-derived spectrum of a vehicle to find the same vehicle in a subsequent image with a false alarm rate of approximately 3×10^{-5} (six out of 200,000) when contextual information providing building and other fixed object locations is used. However, when this type of analysis was attempted for other vehicles in the scene (other blue cars, red cars, etc.), false alarm rates significantly higher were observed. Thus, the feasibility of this method has been demonstrated only for a limited situation.

6. MODEL-BASED ANALYSIS

While the analysis of the MISI imagery is critical to demonstrate the feasibility of vehicle tracking, it represents but a small portion of the potential trade space of situations and sensor capabilities. This is where modeling tools can play an important role in exploring the tradeoffs and parameter sensitivities to enhance our understanding of the capabilities of the technology.

6.1. Model description

The model used to explore sensitivities in this task is the Forecasting and Analysis of Spectroradiometric System Performance[†] (FASSP) model³. This model uses statistical descriptions of targets and backgrounds in a scene and propagates those descriptions through the effects of the remote sensing process to ultimately predict target detection performance.

6.2. Analysis scenario

The model-based analysis scenario was set-up to be similar, although not exact, to the RIT campus environment. Reflectance statistics for five vehicles were estimated from the ground truth measurements. Statistics for the campus-type backgrounds were selected from the FASSP library derived from atmospherically-compensated HYDICE⁴ data collected over a separate, but similar, urban environment. An approximate model for the MISI sensor was developed and used along with the HYDICE model already present in FASSP. Table 2 summarizes the scenario.

Table 2. FASSP model analysis scenario.

Parameter	Value(s)
Target	Blue car 1 (other cars studied as well)
Backgrounds	25% roadway, 15% grass, 15% trees, 15% roof 1, 10% roof 2, 10% bare ground, 5% water, 1% blue car 2, 1% green car 1, 1% green car 2, 1% white car
Visibility	10 km with urban aerosol model
Solar zenith angle	30°
Atmospheric model	Summer mid-latitude
Sensor altitude	1.2 – 5.0 km
Sensor	MISI, HYDICE
Atmospheric compensation	ELM
Detection algorithm	Matched filter

Note that while one car was used as a target for a given analysis, the other four were placed in the background to serve as potential false alarm sources. This models the situation we are studying where we are trying to detect the presence of a particular car in a cluttered urban environment.

The model studies examined the target detection performance for a single hypercube and did not address the “tracking” aspect. For now, the analyses assume the vehicles are moving a significant distance between frames that the problem comes down to one of detection in each given image rather than kinematic tracking.

[†] FASSP was initially developed at MIT Lincoln Laboratory and has been licensed for use at RIT.

6.3. Model sensitivity studies

The first study was to look at the sensitivity to target vehicle color and type using the model MISI sensor. Three vehicles were considered one at a time, with the others placed in the background at 1% of the scene. Figure 8 shows the receiver operating characteristic (ROC) curves for two blue and one green. The “blue car 1” type was easily detectable down to the minimum P_{FA} studied with a $P_D = 1$. The “blue car 2” was modestly detectable at $P_{FA} = 10^{-4}$ and above, while the “green car 1” was not detectable at these low P_{FA} 's. The “blue car 1” was the car easily detected in the same real MISI image as shown in section 5.2. Thus, the model confirms this easy detection and shows that it is not always the case depending on the vehicle paint.

The next study investigated the detection sensitivity to spectral misregistration between visible and near infrared bands. MISI uses separate optical fibers at the focal plane to feed the visible and the near infrared spectrometers. A spatial misregistration has been observed between the 35 visible bands and the 35 near infrared bands. While not measured exactly, the overlap between these bands has been estimated to be lower than 50% for the same pixel index. Figure 9 shows the detection performance for the “blue car 1” for different amounts of pixel overlap between the visible and near infrared bands. As can be seen, as the registration falls to between 25% and 50% the detection rate falls dramatically at the lower false alarm rates. (Note that 25% is the lowest possible overlap; otherwise the misregistration is less in an adjacent pixel). This figure quantifies the impact of one of the MISI artifacts and demonstrates that the results from the empirical analysis may be significantly impacted by this misregistration.

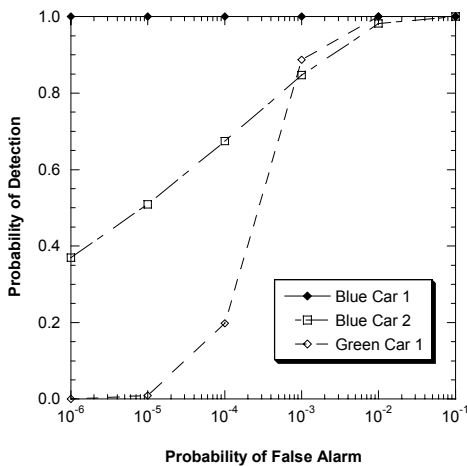


Figure 8. Sensitivity to target car color.

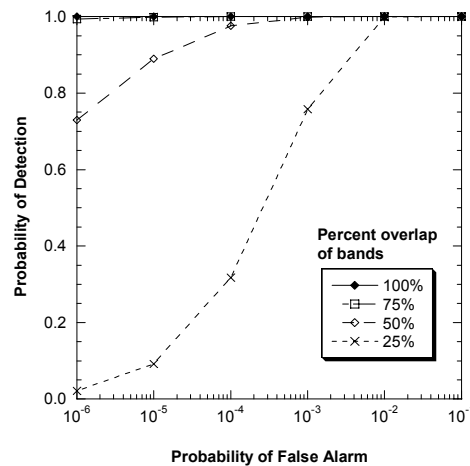


Figure 9. Sensitivity to VIS and NIR misregistration.

The model results presented above assumed the target vehicle occupied at least 100% of a pixel. We know in the case of the MISI data, this was not always the case, so we investigated the sensitivity of detection for the case where the target vehicle was subpixel. Figure 10 shows the sensitivity to the target pixel fill area percentage for the two types of blue cars considered, again using the MISI sensor model (assuming perfect visible and near infrared spectral band registration.) The plot shows probability of detection at a specified false alarm rate of 10^{-5} . As can be seen, the “blue car 1” remains detectable down to about 60% fill factor, while the “blue care 2” type falls very quickly from being only modestly detectable even at 100% pixel fill.

The final study reported compares the performance of the MISI model to two versions of a model HYDICE sensor. HYDICE is a high SNR instrument containing 210 spectral channels spanning 400 to 2500 nm. We considered the problem of detecting a “blue car 2” type of vehicle. Results using a subset of channels (144) corresponding to the atmospheric window regions of the full HYDICE coverage (VNIR/SWIR) were compared to using only HYDICE channels only from the VNIR region (67) and the use of the full channel set (70) from MISI. Figure 11 presents the results which show the dramatic improvement in detection at low target fill percentages from using the additional SWIR channels. These results also demonstrate the benefits of the higher SNR (approximately an order of magnitude) achieved by the model HYDICE sensor compared to MISI, even when using channels covering the same spectral region.

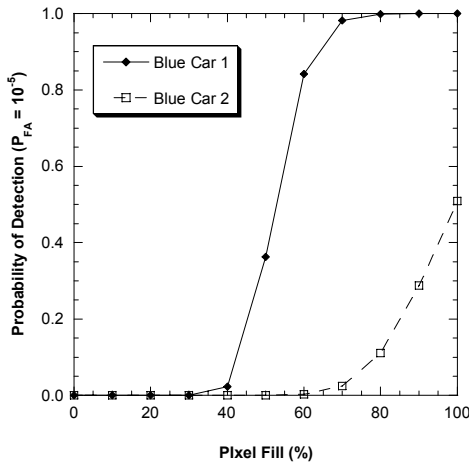


Figure 10. Sensitivity to target car color and pixel fill.

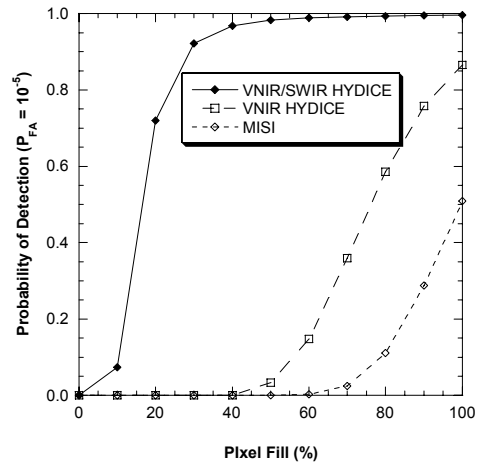


Figure 11. Sensitivity to SNR and spectral coverage (blue car 2).

7. CONCLUSIONS AND FUTURE WORK

The empirical and model analysis results demonstrate the feasibility of using spectral information to uniquely detect a specific vehicle in a given situation. However, the feasibility depends very much on the spectral characteristics of the target vehicle, the complexity of background and the characteristics of the spectral imaging sensor.

While the experiment conducted with the MISI sensor led to a successful demonstration, the relatively low spatial resolution, relatively low SNR, and the visible and near infrared misregistration limited the data quality and our ability to draw additional conclusions from the analyses. The model results supported the empirical findings and were able to explain and enhance our understanding.

Future work will further investigate the data collected by looking at noise reduction pre-processing techniques and the application of more sophisticated detection algorithms. Additional model analyses will be conducted to verify our empirical findings and further explore the parameter trade space. Also, we are considering an opportunity to repeat the experiment with a sensor capable of higher spatial resolution, SNR and expanded spectral coverage.

ACKNOWLEDGMENT

Acknowledgment is made to the numerous additional RIT students and staff who contributed significantly to the experiment and data collection process.

This material is based on research sponsored by AFRL/SNAT under agreement number FA8650-04-1-1717 (BAA 04-03-SNK Amendment 3). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of AFRL/SNAT or the U.S. Government.

REFERENCES

1. D. Messinger and M. Richardson, "Analysis of a multi-temporal hyperspectral dataset over a common target scene," *Proceedings of Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XII*, SPIE Vol. 6233 (this proceedings), 2006.
2. J. Schott, T. Gallagher, B. Nordgren, L. Sanders, and J. Barsi, "Radiometric calibration procedures and performance for the Modular Imaging Spectrometer Instrument (MISI)," in *Proceedings of the Earth International Airborne Remote Sensing Conference*, ERIM, Ann Arbor, MI, 1999.

3. J. Kerekes and J. Baum, "Spectral Imaging System Analytical Model for Subpixel Object Detection," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 40, no. 5, pp. 1088-1101, May 2002.
4. L. Rickard, R. Basedow, E. Zalewski, P. Silverglate, and M. Landers, "HYDICE: An Airborne System for Hyperspectral Imaging," *Proceedings of Imaging Spectrometry of the Terrestrial Environment*, SPIE Vol. 1937, pp. 173-179, 1993.