10-15-2004

High-Resolution Slant-Angle Scene Generation and Validation of Concealed Targets in DIRSIG

Kris E. Barcomb

U.S. Air Force

John Schott

Rochester Institute of Technology

Scott D. Brown

Rochester Institute of Technology

Timothy J. Hattenberger

Rochester Institute of Technology

Follow this and additional works at: <http://scholarworks.rit.edu/article>

Recommended Citation

Kris E. Barcomb, John R. Schott, Scott D. Brown, Timothy J. Hattenberger, "High-resolution slant-angle scene generation and validation of concealed targets in DIRSIG", Proc. SPIE 5546, Imaging Spectrometry X, (15 October 2004); doi: 10.1117/12.561256; <https://doi.org/10.1117/12.561256>

This Article is brought to you for free and open access by RIT Scholar Works. It has been accepted for inclusion in Articles by an authorized administrator of RIT Scholar Works. For more information, please contact ritscholarworks@rit.edu.

High-Resolution, Slant-Angle Scene Generation and Validation of Concealed Targets in DIRSIG

Captain Kris Barcomb^a, Dr. John Schott^b, Scott Brown^b, and Tim Hattenberger^b

^aUnited States Air Force*

^bRochester Institute of Technology, Center For Imaging Science

Digital Imaging and Remote Sensing Laboratory

54 Lomb Memorial Drive, Rochester, NY USA 14623-5604

ABSTRACT

Traditionally, synthetic imagery has been constructed to simulate images captured with low resolution, nadir-viewing sensors. Advances in sensor design have driven a need to simulate scenes not only at higher resolutions but also from oblique view angles. The primary efforts of this research include: real image capture, scene construction and modeling, and validation of the synthetic imagery in the reflective portion of the spectrum. High resolution imagery was collected of an area named MicroScene at the Rochester Institute of Technology using the Chester F. Carlson Center for Imaging Science's MISI and WASP sensors using an oblique view angle. Three Humvees, the primary targets, were placed in the scene under three different levels of concealment. Following the collection, a synthetic replica of the scene was constructed and then rendered with the Digital Imaging and Remote Sensing Image Generation (DIRSIG) model configured to recreate the scene both spatially and spectrally based on actual sensor characteristics. Finally, a validation of the synthetic imagery against the real images of MicroScene was accomplished using a combination of qualitative analysis, Gaussian maximum likelihood classification, and the RX algorithm. The model was updated following each validation using a cyclical development approach. The purpose of this research is to provide a level of confidence in the synthetic imagery produced by DIRSIG so that it can be used to train and develop algorithms for real world concealed target detection.

Keywords: DIRSIG, concealed target detection, hyperspectral image simulation, oblique view, verification and validation

1. INTRODUCTION

Synthetic imagery is a critical component of Imaging Science for many reasons. It allows sensor designers the opportunity to create virtual versions of a sensor without many of the problems associated with creating costly physical versions. Synthetic image generation (SIG) tools also allow system users the ability to determine the best way to utilize a particular sensor design. Those users can model prospective scenes and determine the combination of parameters, such as acquisition time, view angle, and weather conditions that maximizes image quality. Another major benefit of computer modeling is the ability to do detailed error analysis. SIG tools allow a user total control over the process, so that physically impossible experiments, like completely removing the atmosphere or removing noise from the system, are possible. With these tools designers can study exactly where the problems are in the image chain and determine how much each piece contributes to the overall system error.¹ The sensor design, atmospheric conditions, resolution, spectral regions, and targets can all be modified with a fraction of the cost that would be required when acquiring the information from real world scenes.

The main focus of this research is on exploring the utility of synthetic imagery in the area of automatic target recognition (ATR) algorithms. Once a virtual scene is constructed, the parameters that effect the detection capabilities of an ATR algorithm, such as resolution, view angle, and spectral region, can be changed much more readily in a computer than having to create physical versions of that same scene. The algorithms can then be tested against a much more rigorous data set. In addition to ATR validation, SIG tools can potentially train ATR algorithms. Many common ATR algorithms use statistical or non-parametric classifiers and therefore require a

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

great deal of training data. Multiple scenes must be imaged for use as training data for ATR algorithms in order to increase the algorithm's understanding of the different conditions in which a target may be found. This is especially true when training hyperspectral algorithms. Most ATR algorithms suffer from a lack of training data and that problem is compounded when the dimensionality (e.g. number of spectral bands) of the problem is large.² SIG can potentially be a very useful tool for populating the training data when real data is not available. For matched filter or anomaly detection algorithms that do not require training sets, realistic variability within the model becomes key. Too little variability leads to artificially high detection rates.

All of this requires that the SIG model be physically accurate at all levels of the scene that are of importance to the sensor being modeled or the algorithm being trained. The creation of a SIG scene is a very time consuming process because of all of the real world information that must be acquired and cataloged before the simulation can be run. Also, all of the significant underlying phenomenology of the physical world must be understood.

This paper will provide an overview of an effort to validate the Rochester Institute of Technology's (RIT) Digital Imaging and Remote Sensing Image Generation (DIRSIG) tool for target detection algorithms. In the course of the research, truth imagery of the scene was acquired using RIT's Multispectral Imaging Spectrometer Instrument (MISI), a virtual replica scene was constructed, and a detailed comparison of the resulting imagery was conducted. The model is referred to as MicroScene, which gets its name because it is smaller in total area, but is modeled at higher resolution than its counterpart, MegaScene.³

2. DATA COLLECTION

In order to understand how well the synthetic model performs at simulating real phenomenology, real scenes need to be imaged for comparison. This section provides an overview of the truth image collection that was conducted of the MicroScene area at RIT. The imagery was obtained using two imaging instruments owned and operated by RIT, MISI and the Wildfire Airborne Sensor Program (WASP). The WASP sensor is primarily a thermal instrument. It was not used in the validation of the DIRSIG model, but its high resolution, panchromatic, framing array camera was used for fine tuning the spatial locations of objects in the virtual scene.

MISI is an airborne, line scanning instrument with a 6" rotating mirror coupled with a f/3.3 Cassegrainian telescope. The instrument contains many spectral bands. MISI's broadband capability measures the visible, SWIR, MWIR, and LWIR regions of the spectrum. Two separate 36-channel spectrometers cover the electromagnetic spectrum from $0.44\mu\text{m}$ to $1.02\mu\text{m}$ at $.01\mu\text{m}$ increments. The system has been used at RIT for high-altitude aircraft and satellite sensor performance evaluation, data collection for algorithm development, and as a survey instrument for demonstrating proof-of-concept studies in areas ranging from, water quality assessment to energy conservation.⁴

Figure 1 provides a complete overview of the MicroScene area and the locations of the sensors and targets. The sensors were placed on top of a scissor cart in the scene at the location depicted in Figure 1. Then, the scissor cart was raised to an elevation of 50 feet. MISI required a mechanical table to rotate the sensor across the scene to simulate aircraft movement in the along-scan direction. The base of the cart was approximately 100 feet from the center of the target locations which resulted in a nominal resolution of approximately 3 inches.

The primary targets in the scene are three military Humvees. The vehicles were placed at three locations in the scene under various levels of concealment. The first placement was at the box labelled as "Uncovered Humvee" in Figure 1. This was done so that no trees were in front the vehicle and only limited tree cover behind. The second vehicle was located at the box labelled "Humvee In Trees" so that the vehicle was surrounded by trees and only partially visible to the sensors. Finally, the last Humvee was placed at the location of the box labelled "Camo Humvee", where it was draped with woodland camouflage netting. Supplemental camouflage (i.e. twigs, grass, leaves, etc.) were not used for the sake of simplicity. The camouflage draped Humvee was placed in this configuration in the open so that a clear view could be obtained of the camouflage netting, its contours and also the shadowing created by the camouflage pattern. The three Humvee placements can be seen in Figure 2.

In addition to the imagery, many reflectance curves of the surrounding background were collected to populate DIRSIG's material properties database. This was accomplished in the field with an ASD Spectroradiometer and in the laboratory with a CARY 500. Many atmospheric parameters were also measured during the collection.

Figure 1. Sensor and target locations in MicroScene.

Figure 2. The Humvee target under three levels of concealment.

Imagery was collected every hour, on the hour. The full collection began Wednesday the 28th of August, 2003 at 1000. 30 collections were obtained through Thursday at 1500. An image acquired at 1000 on the 29th was selected to validate the model over the visible wavelengths from 400nm to 700nm. The DIRSIG imagery is simulated over the same wavelengths under the same conditions. The virtual MicroScene model will be discussed in the following section.

3. SYNTHETIC SCENE CONSTRUCTION

3.1. MicroScene Model Creation

The DIRSIG model was chosen as the SIG tool for this research. The DIRSIG model is an integrated collection of independent first principles based submodels which work in conjunction to produce radiance field images with high radiometric fidelity. This modular design creates a high degree of flexibility and interchangeability within the model, as well as the capability to diagnose and improve the model by isolating and analyzing each submodel. DIRSIG has evolved over nearly two decades, from a Long-Wave Infrared (LWIR) SIG modeling tool originally developed for ATR algorithm development and sensor trade studies, to a high fidelity model aimed at producing hyperspectral images over the visible through long-wave infrared spectral range.⁵

Figure 3. Various detailed models built for the virtual MicroScene.

Image generation in DIRSIG begins by modeling the geometric information of the scene. 3-D faceted models of the various objects in the MicroScene area were built using the Rhinoceros CAD software package. Figure 3 shows a few of the more detailed CAD models that were built for this scene. From the top-left going clockwise they are: a downwelled radiometer, a shed, a toy wagon with an electric blackbody on it, a generator, a portable weather station, and finally a Humvee. That figure provides a sense of the level of detail that is trying to be attained in this model. These objects were placed in the scene using a combination of the Bulldozer scene placement tool and empirical measurements taken during the actual collection.

The terrain model for the scene is based on a survey of the land that was converted to a gray scale image file. The gray scale values were scaled to corresponded to elevations in meters. Then, a utility was used that converted this image file to the faceted terrain model. Normally, individual facets are assigned unique material properties. This method is insufficient for capturing the variability of the ground in the area. One of the benefits of the DIRSIG model is its hierarchical mapping structure. Once the faceted terrain model was generated it was overlaid with three different mapping images that induce the necessary variability. This is important because it is done without the need for increasing the number of facets used to define the geometry of the terrain or varying the material properties of the terrain on a facet-by-facet basis. The final terrain model combines material, texture and bump maps. The material map is used to distinguish different material types within an object. It is primarily used to create the transition between the grassy regions to the dirt area in front of the shed. The texture map is used to enhance the variability within an individual material type. The one used here is a 3-inch GSD overhead image of the area that has been registered to the terrain model. Finally, the bump map is an image that is used by DIRSIG to characterize the amount of deflection that should be added to an incident ray that impacts the flat facet surfaces. The bump map adds variability to the spectra and also provides the appearance of roughness. The bump mapping effect was used on both the grass and dirt regions.

This hierarchical mapping structure was also instrumental for developing realistic virtual camouflage netting. A three color material map was created from a thresholded digital camera image of the net. The image was thresholded to a level that differentiated the net from the holes in the net. The white areas were recolored to

Figure 4. Diagram of DIRSIG's treatment of NULL material mappings.

Figure 5. Threshold generated material map for camouflage netting.

(a) Digital camera image of the Humvee underneath camouflage netting. (b) DIRSIG image of the Humvee underneath camouflage netting.

Figure 6. Truth vs. DIRSIG imagery comparison of a Humvee underneath camouflage netting.

represent the pattern of the net's two different camouflage colors. Rather than point to a measured material file, the black areas were assigned a NULL material ID in the lookup table portion of the DIRSIG configuration file. When DIRSIG casts a ray that hits the portion of a material mapped facet with a NULL material ID, the ray is allowed to pass through the facet as if it wasn't there. A graphical representation of the how DIRSIG treats NULL material mappings is shown in Figure 4. The capability reduces the modeling needs that would otherwise be required if each hole in the net was cut from the model on a facet basis. The camouflage material map is shown in Figure 5. The black areas in Figure 5 correspond to the areas assigned with the NULL ID tag in the material lookup table.

The dramatic results of this process are shown in Figure 6. The image on the left was taken with a standard digital camera from underneath the camouflage net around midday. The image on the right was created with DIRSIG by placing a similar "synthetic" camera under the synthetic net. The intricate shadow pattern on the vehicle is apparent in both images. It should be noted that DIRSIG's BRDF model, which helps to determine a more realistic background shape factor for each facet, was turned on for this image. The run time was dramatically increased, but the result was a much more realistic image. This side-by-side comparison is an example of the level of detail that can be achieved in DIRSIG.

Creating the surrounding vegetation was the last major portion of generating the spatial qualities of the virtual scene. There are 4 different types of tree models in the synthetic scene that were all produced using the TreePro vegetation modeling software. Each tree model is loaded into memory only once regardless of the number of times it is instantiated in the scene. So, each one is scaled and rotated in various ways to produce

Figure 7. Spatial comparison of WASP and DIRSIG overhead imagery.

variability without running into system resource problems. A simple, but powerful utility was created for placing the trees in the scene quickly and accurately. The software tool, called TreePlanter, converts points on an overhead image with a known GSD into DIRSIG formatted text files that contain the point locations in DIRSIG scene coordinates. The images in Figure 7 shows how accurate the technique was in recreating the forest region around MicroScene. The image on the left was taken with the panchromatic imager on WASP and the image on the right was rendered in DIRSIG. Spectral considerations are not important for comparing these two images. The WASP image was taken in the fall after the leaves had changed colors causing some trees to appear much brighter than they did in the summer. They are meant to show that if overhead imagery is available, then realistic forests can be populated with a fraction of the effort that has been put forward in the past.

3.2. Synthetic Image Post-Processing

The information necessary to characterize the geometric effects of MISI (focal length, detector size, scan rate, and duty cycle) were all known prior to this research. Also, the spectral response characteristics were known because the gains, biases, and spectral band-passes of the system had recently been evaluated. The PSF and noise characteristics were still unknown, though, and needed to be determined for use in post-processing the DIRSIG version 3 generated imagery[†].

Insufficient resources were available to do extensive laboratory testing of MISI's PSF. So, it was approximated from the truth imagery by measuring a slice of the imagery across a high contrast transition region. A vertical transition between the bright sky and a dark building were used here. The sky-building transition acts like a knife edge. The derivative with respect to pixel location across the transition can be used to estimate the PSF in one dimension. Since no significant difference was observed when a horizontal slice was taken across the side of the building, a circularly symmetric gaussian function was used based on the one-dimensional slice. A 7x7 PSF kernel was generated by fitting a continuous Gaussian function to the data points and then resampling that function so the value of the kernel's center pixel matched the area normalized Gauss' peak value. The DIRSIG imagery was convolved with this kernel to simulate the effects of MISI's optical system.

The final post-processing step in creating sensor specific SIG imagery is to add the system noise. When MISI takes images, the instrument is configured to gather 40-50 lines of data while the shutter is closed so that a dark region shows up at the beginning and end of each image. After the bias is subtracted out of the images the only information left is a result of system noise, which is assumed to be additive and have a zero mean value. The standard deviation and band-to-band correlation of MISI's system noise were measured from these

[†]DIRSIG 4, which is currently in development, will incorporate sensor noise characteristics into the resulting SIG imagery during the rendering process

dark scan regions. To generate synthetic images with the same noise characteristics, these dark regions needed to be synthesized into noise cubes with the same spatial dimensions as the images being generated by DIRSIG. This was accomplished through a principle component (PC) transform of just the noise region of the calibrated MISI imagery. The transform serves to entirely decorrelate the noise. Then, the standard deviations of each transformed noise band were used to generate an uncorrelated noise cube of the same spatial dimensions as the DIRSIG imagery using a simple random number generation routine. Next, the synthetic noise cube was inverse transformed using the same covariance matrix generated by the original forward PC transform. Finally, the resulting noise cube was simply added to the DIRSIG imagery. Figure 8 shows a graphical representation of this process.

Figure 8. Process flow for creating correlated noise for DIRSIG imagery.

4. VALIDATION

This research project is an effort to understand how well DIRSIG phenomenologically simulates the real world at high spatial resolutions for the purpose of ATR algorithm training. Three primary methods were examined for evaluating the success of the MicroScene model in achieving that goal. This section will begin with a qualitative comparison of the imagery and the spectra of some of the key targets and areas in the imagery. This method of comparison was the most useful for determining necessary improvements. It will be used continuously throughout the on going cyclical development of the MicroScene model.

Once the scene geometry and radiance spectra were sufficiently accurate, two forms of target recognition algorithms were used on the data. First, a Gaussian Maximum Likelihood (GML) classifier was used to classify the imagery. Second, specific target detection analysis was accomplished using the RX algorithm.^{6,7} It is important to recognize that this effort attempted to generate a synthetic scene that had the same features as the real scene. This means that the goal was for selected targets to be as close as possible and for the synthetic backgrounds to be within the range of actual backgrounds, but not reproduce them at the element by element level.

4.1. Qualitative Image Comparison

Figure 9 shows the truth image on the left and DIRSIG image on the right. The spatial quality of the scenes are very similar. Remember, everything in the scene is empirically derived. The locations of all of the objects were taken from overhead imagery or were measured during the collection. MISI's focal length, duty cycle, PSF, noise, and detector properties were all put into the model as is and the resulting synthetic imagery is very encouraging. The imagery shows that DIRSIG is capable of capturing the spatial distortions and tangential effects of very unique sensor configurations like the one used in this experiment.

The only issue that has been found with modeling high resolution scenes at this kind of oblique viewing angle can be seen in the background of the DIRSIG imagery. The dirt region, that is meant to only be in front of the shed, gets repeated off in the distance. This occurs because DIRSIG tiles the material maps, described in Section 3.1, when the object is larger than the map. The material map used for the terrain needed to be extremely large to create the necessary detail in the grass-to-dirt transition regions, but it would have been too large to fit into memory if it would have been created to fit the entire terrain map at this resolution. To the eye, this issue is barely noticeable, but it will come back up later when the anomaly detection algorithm is run on the image. If required, this limitation can be overcome by either substituting lower resolution maps or adding more memory to the systems running the simulation.

Figure 9. MISI vs. DIRSIG image comparison.

Figure 10. MISI vs. DIRSIG spectral comparison in the visible region.

Finally, the blurry portion of the bottom of the MISI imagery is a result of the line scanner going out of focus as the scan mirror moves the field of view closer to the sensor. Modeling this effect was not a goal of the research and so the lower quarter of the MISI imagery should be disregarded for all of the analysis in this paper.

4.2. Qualitative Spectral Comparison

The first portion of the development cycle for the MicroScene model was primarily focused on recreating the geometric properties of the scene. Once the spatial locations, shapes, and sizes of the models were accurate, the next priority was getting the spectral qualities of the model correct. The spectral radiance accuracy for the visible region of the spectrum is shown in Figure 10. The spectra in the figure were obtained by averaging the pixels in a region of interest over each of the targets.

The spectral accuracy of most of the important target locations in the scene are very good. Ironically, the control targets used in the scene were some of the most difficult to model spectrally. The black tarp in Figure 10 is representative of all of those targets. The radiance of the control targets in the MISI imagery is much higher than in the DIRSIG simulations. It is believed that this is largely a result of the sun-target-sensor viewing geometry. At 1000, the sun was almost directly in front of the sensor which created a near perfect specular angle between the sensor, the control targets and the sun. As of this paper, bi-directional reflectance effects have not been included, so materials with significant non-lambertian qualities will not be modeled accurately.

The spectral accuracy of the DIRSIG imagery is generally very encouraging. The next step in the development cycle was to see if the model was capable of training and developing classification and detection algorithms.

Figure 11. MISI vs. DIRSIG GML classification comparison.

4.3. GML Classification Comparison

After significant data collection, scene development, and sensor characterization, the model was ready to be tested for its primary purpose of training and developing detection algorithms. Detection algorithms generally look for small targets, but the modeler cannot focus solely on the accuracy of those few pixels that encompass the target; the entire synthetic image must be accurate. It is very important that a diverse background surrounds that target so that the algorithm will have sufficient variability to perform as it would on real data. Therefore, the Gaussian maximum likelihood⁸ classification algorithm was selected to determine if data derived from the MicroScene model could be used to populate a training set for classifying the truth imagery. Six target classes, the dirt, shingles, Humvee, trees, grass and white tarp, were chosen from the DIRSIG imagery. The Humvee class was taken from the Humvee in the open. The resulting classification image is shown in Figure 11(a). For comparison, the same training set was also used to classify the DIRSIG image that it was derived from. The results of that classification are shown in Figure 11(b).

The MISI imagery still displays more variability than the DIRSIG imagery in those two figures, but the classes line up very well between the two data sets. The classifier found all of the trees, the portion of the dirt that was in direct sunlight, and most of the grass. Portions of the Humvee are classified correctly, but it is not surprising that the Humvee was also classified as dirt, since the Humvees in the scene had recently been through training exercises and had not been washed. Portions of the shingles were also classified correctly while the rest of the shingles was classified in the Humvee class. It is interesting to note that this occurs even in the DIRSIG model where the shingles are in the shade of the trees. The Humvee class also shows up in the trees, grass, in the shadows, and around the camouflage netting. This is not all that surprising either, since the Humvee camouflage paint is meant to blend in with these objects. Also of note is that the pixels around the outsides of the trees in both classification images get assigned to the dirt class in many places. This suggests that DIRSIG's material mixing algorithms are working properly to create realistic transition regions between materials.

The classification image in Figure 11(c) was made by using training data of the same classes, but this time they were derived directly from the MISI imagery. There are only a few differences between this image and the DIRSIG derived classification in Figure 11(a). First, the white tarp is classified correctly. The DIRSIG derived white tarp class was unable to do this probably because the tarp is too non-lambertian. Second, the buildings in the background are classified as trees. This happens in the DIRSIG imagery, but it does not occur in the MISI image classified with DIRSIG data. Finally, the Humvee class only shows up in the trees in a few pixels, not in the abundance that it did in Figure 11(a). Overall, though, the DIRSIG derived training data does almost as good a job at classifying the MISI imagery as the truth derived training data.

Also of note are the unclassified portions of the imagery, which correlate well in all three images. The most obvious region is the sky, which was not an input class, but also many of the small targets in the scene are not classified. The classifier leaving the small targets out of the general classification suggests that the anomaly detection algorithms will be able to exploit those targets. The next section will explore this idea in more detail through the RX algorithm.

4.4. RX algorithm results

The constant false alarm rate (CFAR) version of the RX algorithm was chosen for this validation. RX is primarily used for detecting small targets and can be used in either anomaly detection or matched filter modes. This algorithm relies on the assumption that the image clutter can be described as a Gaussian random process with a fast spatially varying mean and a more slowly varying covariance. The RX algorithm uses a combination of spatial and spectral information to detect targets in an image through a convolution-like process. Essentially, it works by comparing a spatial subset of a multi- or hyperspectral image to its surrounding neighborhood and then producing a scalar result based on how different the subset is from that neighborhood. This is done across the entire image and the finished product of the algorithm is a grayscale image map that can be thresholded to identify potential targets in the imagery.

RX performs well in regions where the targets are small, relative to the neighborhood size chosen, and are significantly different from their background. While an exhaustive analysis of the algorithm was not accomplished, it was determined that the best results were found when a 45x45 pixel neighborhood was used with a 5x5 spatial subset. The success criteria for this research is not to determine the best way to use any particular algorithm, but to generate synthetic imagery that produces similar results to real imagery under target detection analysis.

The results of running the RX algorithm in anomaly detection mode on the MISI and DIRSIG images are shown in Figure 12. The gray scale image produced by the algorithm was thresholded and then the entire image was dilated with a 9x9 kernel for presentability. In that figure the circled regions represent known small targets in the scene, such as the generator and the blackbody, that were detected in both images.

Most all of these targets were detected in both images. Although, the Humvee in the trees was not detected in the DIRSIG imagery. A comparison of that Humvee in both images shows that its radiance spectra is quite a bit less in the DIRSIG imagery. Therefore, it stands out from the tree less than it does in the truth imagery. Since the material file used for the Humvees is accurate, based on the comparison shown in Figure 10(e), the problem must lie in the vehicles placement beneath the trees. A small shift in spatial location may fix the problem. All other points shown in the anomaly detection images are either false alarms or interesting targets that weren't modeled. For example, one of the targets is over a grouping of small white flags that were placed in the grass for another experiment that was going on during the collect. These flags, while not expected to interfere with the image collection, proved to be significant to the detection algorithm.

Some of the false alarms in the DIRSIG imagery are a result of the material map tiling issue discussed earlier. The dirt region repeats in the distance. As it gets farther out, its total size in pixels gets closer to the 5x5 square target region of the RX kernel. The algorithm only sees dirt in the target pixels and grass in the neighborhood pixels. More time spent modeling detail into the background would help decrease these artificial false alarms.

Figure 13 shows the results of the RX algorithm running on both sets of data in matched filter mode. The input spectra for the algorithm was selected as the mean radiance spectra of the Humvee in the open. That seemed like a reasonable amount of a priori knowledge for running this kind of detection in the field. The threshold point for the RX results was chosen at the point where the Humvee in the trees was just classified as

Figure 12. MSI vs. DIRSIG imagery comparison after RX anomaly detection.

Figure 13. MSI vs. DIRSIG imagery comparison after RX matched filter detection using DIRSIG derived spectra of the Humvee in the open.

a target in the truth imagery. The same point was used in the DIRSIG image. Target points in the two images in Figure 13 are only circled if they are over the location of one of the Humvees.

This time the algorithm finds the Humvee in the trees and also the Humvee in the open in both images. In the DIRSIG image, the Humvee under the camouflage netting also gets detected. The truth imagery contains a quite a few other, seemingly random, false alarms scattered across the image. Analysis of the truth imagery in these locations did not reveal any obvious reason for this. While not as prevalent in the DIRSIG image, a few of the false alarms in that image are also over regions where there is no target of interest.

In all, the algorithms produced similar results. Although there are identifiable differences, many of them can be attributed to spatial detail that could be fixed with more time spent in the geometric modeling process. Also, the algorithm identified objects in the truth imagery that were thought to be inconsequential at the time the imagery was collected, but proved to be significant from the algorithm's perspective. Where these objects are identifiable, as is the case with many of the points in the anomaly detection imagery, these objects can be added to the model. This shows that these algorithms can not only validate, but also help to identify improvements and provide direction into the next step in the continuous improvement of the model.

5. CONCLUSION

This paper provided an overview of an effort to validate DIRSIG's ability to model high-resolution, slant-angle scenes for use with target detection algorithms. Overall the results were very encouraging. A qualitative comparison of truth and synthetic imagery showed that DIRSIG can recreate the spatial effects of unique sensor configurations. It also has the ability to render highly detailed models and intricate shadow patterns. Mean radiance values of a number of objects in the scene were compared and the model also performed extremely

well for all near-Lambertian materials. This was expected because bi-directional material properties were not available at this point for inclusion in the model.

The spectral variability of those materials was also examined through the use of GML classification and the RX algorithm. The truth imagery was classified with both DIRSIG and truth derived training sets. The two classification results show some differences, but overall, the objects in the image were classified appropriately with either training set. The results of the RX algorithm are also encouraging. Not only did the algorithm help to validate the model's spectral and spatial accuracy, but it also pointed out areas for the model's improvement that can be incorporated into the next development cycle.

REFERENCES

1. J. Schott, *Remote Sensing: The Image Chain Approach*, Oxford University Press, Oxford, NY, 1997.
2. R. Duda, P. Hart, and D. Stork, *Pattern Classification*, John Wiley and Sons, Inc., New York, NY, second ed., 2001.
3. E. Ientilucci and S. Brown, "Advances in wide area hyperspectral image simulation," in *Targets and Backgrounds IX: Characterization and Representation, Proceedings of SPIE* **5075**, pp. 110–121, April 2003.
4. CIS, "<http://www.cis.rit.edu/research/dirs/research/misi.html>." Online, 2003.
5. J. Schott, S. Brown, R. Raqueno, H. Gross, and G. Robinson, "An advanced synthetic image generation model and its application to multi/hyperspectral algorithm development," *Canadian Journal of Remote Sensing* **25**, June 1999.
6. I. Reed and X. Yu, "Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution," in *IEEE Transactions on Acoustics, Speech, and Signal Processing*, **38**, pp. 1760–1770, October 1990.
7. I. Reed, X. Yu, and A. Stocker, "Multidimensional signal processing for electro-optical target detection," in *Signal and data processing of small targets, Proceedings of the SPIE* **1305**, pp. 218–231, 1989.
8. J. Richards, *Remote Sensing Digital Image Analysis*, Springer-Verlag, Berlin, 1999.