

Advances in Wide Area Hyperspectral Image Simulation

Emmett J. Ientilucci and Scott D. Brown

Rochester Institute of Technology, Center for Imaging Science
Digital Imaging and Remote Sensing Laboratory
54 Lomb Memorial Drive, Rochester, NY, USA

ABSTRACT

With the advent of hyperspectral imaging spectrometers comes the need for procedures that detect and interrogate spectral quantities of interest. Such procedures or algorithms play a key role in the dissemination and interpretation of hyperspectral data. Validation of these algorithms involves well-characterized field collection campaigns that can be time and cost prohibitive. Radiometrically, as well as geometrically, correct synthetic imagery offers algorithm developers a surrogate to potentially unattainable field campaigns. The image simulation surrogate must ideally match real world scenes in both spatial and spectral complexity for one to have faith in algorithm performance. To this end, there is a need to develop synthetic scenes, based on real world data, which encompass full 3-dimensional geometric complexities as well as wide-area, spectrally complex backgrounds. Prior work has been done on the inclusion of backgrounds into a synthetic environment, however, this work did not generate wide-area imagery with all the complexities realized in real world data.

This paper investigates the generation of a wide area synthetic scene rendered by the Digital Imaging and Remote Sensing Image Generation (DIRSIG) model. The large area scene or “MegaScene” described in this paper is 0.6 square miles and contains an order of magnitude increase in objects, materials, and spectra, as compared to previously rendered scenes. Hyperspectral analysis using off-the-shelf classification and target detection algorithms was performed on the data to illustrate quantitative and qualitative fidelity.

Keywords: DIRSIG, hyperspectral image simulation, wide area coverage

1. INTRODUCTION

1.1. The DIRSIG Model

The Digital Imaging and Remote Sensing Image Generation (DIRSIG) model is a complex synthetic image generation application which produces simulated imagery in the visible through thermal infrared regions.¹ The model is designed to produce broad-band, multi-spectral and hyperspectral imagery through the integration of a suite of first principles based radiation propagation submodels. These submodels are responsible for tasks ranging from the bi-directional reflectance distribution function (BRDF) predictions of a surface to the dynamic scanning geometry of a line scanning imaging instrument. In addition to submodels that have been specifically created for the DIRSIG model, several of these components (*e.g.* MODTRAN and FASCODE) are the modeling workhorses for the multi- and hyperspectral community.

1.2. Historical Perspective

The maturity of scene simulation tools for use in the evaluation of hyperspectral algorithms has been explored in the past.² The resulting imagery lacked the extensive spatial and spectral complexity present in real world image data and the performance of common hyperspectral algorithms run on the synthetic data was found to be extremely optimistic. It was determined that the simulated imagery needed significant improvements in order to match the spatial and spectral complexity of real world image data.

The scene simulation work being conducted at the Rochester Institute of Technology (RIT) has been focused on developing a radiometrically rigorous scene simulation capability. The DIRSIG model is capable of producing imagery that features mixed pixels, complex in-situ illumination loadings and the spectral statistics observed in real world materials. We were confident that the physics implemented by the model was extensive enough to create realistic imagery when robust input databases were utilized and the proper sensor modeling was employed.

Further information: www.cis.rit.edu/~dirsig
emmett@cis.rit.edu, phone: (585)475-7778
brown@cis.rit.edu, phone: (585)475-7194

1.3. Goals and Scope

One of the primary goals of this modeling effort is to produce imagery that can be used to test the performance of spatial and spectral image exploitation algorithms. If the modeling tools can successfully reproduce imagery with spatial and spectral “clutter” comparable to real-world imagery, then candidate algorithms can be extensively tested over a wider range of environmental conditions at a significant cost savings over field collections. In addition to the convenience of creating synthetic data over a wide variety of environmental conditions and image acquisition approaches, a synthetic data set can provide “perfect” ground truth without the technical and logistical challenges of running a successful field campaign. As we will discuss later, the synthetic image ground truth can provide the algorithm tester with very detailed insights into the performance of the algorithm that cannot be matched by field collections. It should be noted, however, that the authors do not believe that synthetic data should ever wholly replace real-world image data for algorithm testing. Instead, synthetic data should be considered a powerful tool to assist in the testing of algorithms and potentially as a surrogate when real data is not available.

This paper will discuss recent advances made in the wide area hyperspectral image simulation. Specifically, we will discuss the importance of detailed scene construction techniques leveraged by carefully planned and executed ground-truth collections. We will also discuss and demonstrate the effect that modeling fidelity has on algorithm performance. The scope and results discussed in this paper are limited to the solar or reflective region of the spectrum.

2. BACKGROUND

To produce imagery that contains the spatial and spectral complexity of field collected data, the image generation model must be able reproduce a large set of radiative mechanisms that combine to produce the spectral signatures that are collected by real world imaging instruments. The DIRSIG model attempts to incorporate a wide array of these image forming processes within one modeling environment. To drive these predictive codes, the model must have access to robust characterizations of the elements to be modeled. For example, input databases describe everything from a chemical description of the atmosphere as a function of altitude to the spectral covariance of a specific material in the scene. In many cases, access to these supporting databases is the limiting factor in the fidelity of the produced imagery.

Once the properties of elements in the synthetic scene are well defined, the radiative transfer model is responsible for predicting the movement of photons within the scene and eventually into the modeled sensor. In this paper, we will focus on how the DIRSIG model addresses four (4) important aspects of scene modeling: (i) geometric complexity, (ii) spatial-spectral diversity, (iii) directional reflectance and illumination loading and (iv) spectral mixing.

2.1. Geometric Complexity

The DIRSIG model utilizes scenes that are constructed of facetized and functional surfaces. A ray tracer is used to determine surfaces along paths within scene that contribute to radiance fluxes at specific points. In general, scene modeling tools have sought to strike a compromise between scene coverage and geometric detail. Because of the labor resources associated with their creation and the run-time resources required to execute the model, large-area scenes have typically contained elements that are modeled with less geometric detail. Without geometric detail it is difficult to model some of the radiometric mechanisms that arise in objects with significant geometric complexity (*e.g.* a tree canopy). The DIRSIG model has been historically used for modeling small areas with a high level of geometric and radiometric fidelity. DIRSIG scenes that have been created in the past were commonly on the order of fractions of a square kilometer. At the onset of the MegaScene project, the goal was to demonstrate that the model could be used to simulate scenes that were many square kilometers in size without a compromise in either geometric or radiometric fidelity. In addition to large area coverage, the final scene demonstrates a level of radiometric complexity much greater than that seen in past renderings.

The scene geometry for the MegaScene was constructed to mimic an area on the northeast side of Rochester NY. The area has a combination of urban and suburban residential, industrial and forested areas. This area is part of an intensive collection region that RIT regularly images (either with in-house or externally tasked sensors) and collects ground truth within. The proximity to the RIT campus allows for easy access to the candidate area thus allowing for rigorous data collections of surface optical properties.

In this paper, we will discuss the first completed region or “tile” that measures 1.4 square kilometers within a much larger region. The scene construction process began using 10 meter USGS digital elevation model (DEM) data which

was faceted by an algorithm optimized for fitting elevation data with faceted surfaces.³ The remaining elements in the scene were created using a suite of off-the-shelf tools that produce faceted models.

2.1.1. Source Geometry

The residential housing in the area was surveyed and a candidate list of ten (10) geometrically unique housing types was created. A list of specific commercial and government buildings to be constructed was also created. Faceted models of all the buildings were developed in the Rhinoceros⁴ computer aided design (CAD) package (see Figure 1a). In addition to the buildings, a variety of other scene elements including vehicles, swimming pools, etc. were built with Rhinoceros. A survey of the trees in the area yielded a set of six (6) different species that would need to be geometrically modeled. For this task, the Tree Professional software package⁵ was utilized which allows the user to create species specific 3D models of various trees (see Figure 1b). For the MegaScene project, three (3) geometric variants were created for each species.

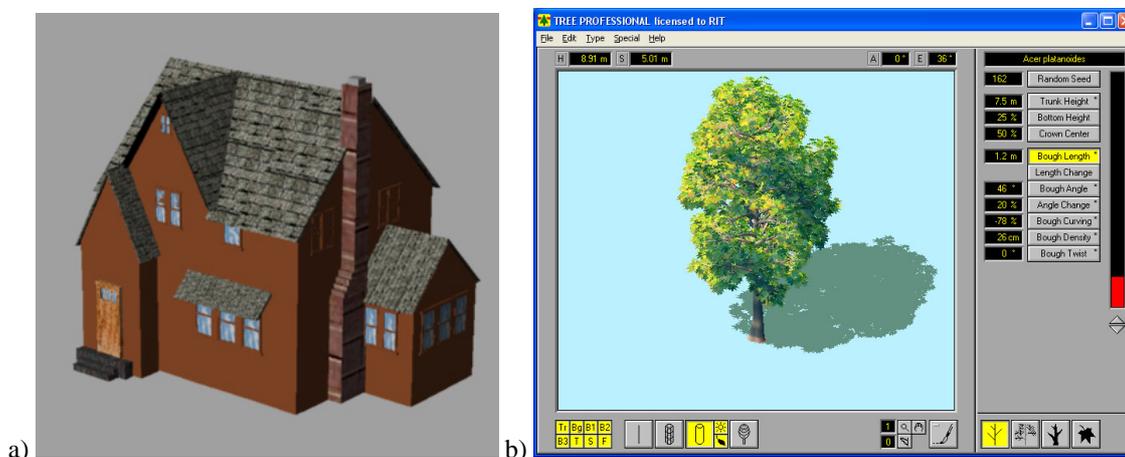


Figure 1. (a) Example of a house generated with Rhinoceros and a (b) screen grab of the Tree Professional software package illustrating a rendered Norway Maple.

2.1.2. Geometric and Attribute Instancing

Instancing is the general practice of using a smaller number of basis objects to simulate a large number objects. *Attribute instancing* is the practice of using a single geometric object to create a larger set of objects composed from the same geometry but with different attributes. For example, a single house style may be used to create a set of houses of the same style but with different roofing and siding materials. *Geometry instancing* entails the use of translation, scale and rotation transforms that are applied on-the-fly to a base object to create a larger number virtual objects that use a fraction of the overhead associated with a full copy of the object. For example, geometric instancing can use a set of exemplar trees consisting of a few hundred thousand facets to model an entire forest containing thousands of trees that would potentially require billions of facets to model.

The MegaScene makes extensive use of both attribute and geometry instancing. The base set of houses were attributed with different materials to produce a set of fifty (50) geometrically and spectrally unique houses. The same approach was applied to the trees to produce a set of twenty (18) geometrically and spectrally unique trees. Objects were manually inserted onto the terrain base using an in-house software tool developed specifically for this task. This graphical insertion tool allows the user to project co-registered air photo data (in this case Kodak CitiPix imagery) onto the terrain so that it can be used as “tracing paper” in order to guide the insertion of buildings, houses, trees, etc. Each geometric instance of these elements in the scene can have a different set of translation, scale and rotation values associated with it. The first tile in the MegaScene contains over 5,000 object instances which account for a virtual facet count of over 500,000,000 facets.

2.2. Spatial-Spectral Diversity

Image simulation models are faced with the challenge of modeling a potentially infinite number of optical elements in the scene that have unique optical properties. The modeling community must invent ways to introduce the spatial and spectral complexity observed in real world scenes using a finite number of inputs.

The strategy used by the DIRSIG model begins with the assumption that a real world scene contains a finite set of classes or endmembers that an image analyst or algorithm would use to describe the scene. These will include common background materials such as asphalt, grass, soil, building materials, etc. in addition to target specific materials. In some cases, this class list may be further subdivided into fresh versus aged asphalt, healthy versus stress grass, etc. These classes arise from the grouping of regions within the image or scene that have similar spatial and/or spectral characteristics. The MegaScene has a top-level material map that defines the various materials used on the terrain surface for regions of the scene. This material map was created using a traditional classification algorithm on CitiPix RGB imagery of the actual site. It was found that the Gaussian Maximum Likelihood (GML) classification algorithm yielded acceptable results for the six (6) primary terrain materials (new asphalt, old asphalt, healthy grass, stressed grass, etc.). The output of the GML was used to attribute the surface terrain. All other materials were introduced via the standard DIRSIG per-facet attribution process. There are currently 110 unique materials within the first tile.

Within each of these materials or classes exists some degree of variation which is commonly referred to as *texture*. The appearance of texture in observed imagery results from spatial variations in reflectance (arising from inhomogeneities and natural variation within the material), orientation, surface structure, shading or a combination thereof. These spatial and spectral patterns are characteristic of the material itself, and they introduce a degree of uncertainty to the analyst and algorithm alike regarding the categorization of a given observation. For example, a region of healthy grass and region of stressed grass may eventually meet within a larger region. That boundary must be defined by some threshold that categorizes one location as “healthy” and another as “stressed” when in fact they may be almost identical.

To simulate texture in targets, DIRSIG utilizes a large database of reflectance curves for a given material (presumed to represent the variations from inhomogeneities)⁶ and a directional reflectance model to introduce variances due to orientation and surface structure.⁷ To introduce spatial variations in spectral reflectance, a texture image is assigned to the material class which represents the spatial variation in reflectance for a specified wavelength region. During the rendering process, a mapping mechanism identifies a pixel in the texture image that is then used to drive the selection of a spectral reflectance curve from a large database of spectral measurements. This mapping technique uses a statistical mechanism that relates the variation in the texture image to variations in the spectral database (see Figure 2). The selected spectral reflectance curve is then mapped to the geometric location and utilized in all the spectral computations involving that location.

The most critical element of this modeling process is the availability of well characterized reflectance spectra for each material. To correctly introduce the spectral variance and covariance observed in real materials, the texture algorithm must have access to a large number of reflectance curves that define the spectral mean, variance and covariance for each material. Historically, spectral reflectance databases have contained limited observations for each material. Many of the databases that contain multiple measurements for a grass material, for example, actually contain measurements of different *types* of grass as opposed to many observations within one region of grass. In order to feed this simulation activity, RIT embarked on an extensive ground collection campaign aimed at building spectral reflectance databases for materials within the MegaScene area. Using an ASD Field Spectrometer, hundreds of materials were measured using a collection scheme where each material was extensively characterized with multiple observations. The complex geometry of the trees, buildings and other objects used within the scene introduce variability due to orientation and shadowing. Later, we will discuss how the mixing process introduces additional structure to the spectral statistics of the scene.

2.3. Directional Reflectance and Illumination Loading

The reflected energy from a surface is a product of the the illumination loadings incident upon the surface and the geometry specific reflectance factor for each one of those loadings. These geometry specific reflectance factors can be described for real world materials by the bi-directional reflectance distribution function (BRDF). Many models make assumptions that the radiance by a surface is simply the directly reflected solar energy. In some cases, models will incorporate a diffuse term that uses a diffuse reflectance in conjunction with an estimate of the diffuse illumination load. In many cases, this diffuse illumination will be based on a unobstructed sky assumption. In reality, the diffuse

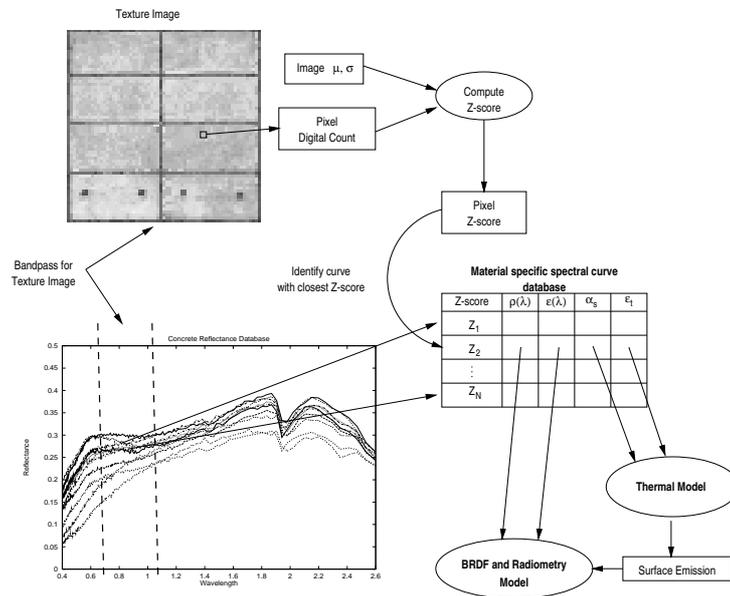


Figure 2. A flow diagram illustrating the texture methodology that is used to introduce spatial variations in spectral reflectance.

illumination load includes indirect solar energy that is scattered onto the surface from nearby objects and the sky. When a target of interest is in shade (the absence of direct solar illumination), the indirect solar load is the major illumination source. When this illumination is scattered by the sky, clouds or trees we refer to this illumination as “sky shine”, “cloud shine” and “tree shine”, respectively. These illumination sources have been given special names because the hyperspectral community has discovered that these sources “color” the target of interest and make it harder to detect using traditional spectral detection methods.

The DIRSIG model attempts to incorporate many of these challenging mechanisms into the predicted signatures by using a directional reflectance model and a background sampling method. In operation, the DIRSIG ray tracer casts rays into the hemisphere above the target to predict both the direct and indirect illumination load incident upon the target. For each of these rays, the incident irradiance is calculated, weighted by the geometry specific reflectance factor and then added to the reflected radiance. This incident irradiance calculation incorporates the contribution from multiply scattered photons because the directly reflected energy from “background” surfaces also is computed. For example, if one of the background rays hits a leaf, the sunlight scattered by the leaf back toward the target will be computed. The same holds true for a cloud, building, etc. As a result of this approach, the DIRSIG model has been used to demonstrate the influence of phenomenological effects including tree shine on the reflected radiance spectrum of a target.

2.4. Spectral Mixing

Spectral mixing is the term that we used to describe the integration of several unique or “pure” spectra within the area of a pixel element on the focal plane. The resulting measured spectrum is a weighted combination of the “pure” elements that were contained within the footprint of the projected pixel on the ground. For example, if the footprint of a pixel overlaps the transition of a roadway into dirt, the resulting pixel spectrum will not represent either the roadway or the dirt, but rather a combination of the two. The mathematics of this “mixing” process is usually categorized into two assumptions: *linear* and *non-linear* mixing. Under the linear mixing assumption a set of spectrally constant weighting factors can be defined that combine the pure elements into the final spectrum. Under non-linear mixing assumptions, these weighting factors are not spectrally constant.

The DIRSIG model produces mixed spectra via two mechanisms. Each pixel is sub-sampled by the sensor model, and the sub-detector spectral radiances are combined using the weights defined by the spectral point spread function (PSF) for the pixel element. This spatial mixing within a pixel element on the focal plane is a linear mixing process. Since the DIRSIG radiative transfer engine models transmissive elements, any path within the scene can also be thought of as a non-linear mixture of a series of transmissive and non-transmissive elements. For example, a path that intersects

a tree leaf and the ground will include the radiance reflected by the leaf and the radiance reflected by the ground under the leaf. However, the ground radiance will be weighted by the transmission of the leaf which is not spectrally constant. As a result, the final spectrum is not the sum of the two participating spectra weighted by two scalar mixing fractions.

3. SIMULATIONS

To evaluate both the qualitative and quantitative fidelity of the scene and simulation system a series of output images were produced. The hyperspectral image cubes were produced using a simplified model of the HYDICE imaging system. The modeled imaging platform was flown straight and level at an altitude yielding a 1 meter ground sampling distance (GSD). The atmosphere for this simulation was the MODTRAN “Mid-Latitude Summer” built-in atmospheric profile without any aerosol contributions. The acquisition time of day was chosen to provide high sun illumination conditions. The final image cube (see Figure 3) measures 320 pixels wide by 1,350 scan lines with 210 spectral channels (as defined by a nominal HYDICE spectral calibration).



Figure 3. (a) HYDICE simulation and three (b-d) additional close-ups of the scene.

In these initial simulations, the sensor was modeled without the spectral “smile” (spatially varying response) and the spectral noise associated with the actual system. Therefore, the algorithm results that will be presented will still reflect optimistic performance compared to the imagery from the actual instrument. A more appropriate atmosphere and the inclusion of these additional instrument characteristics will be incorporated into later simulations and performance studies.

3.1. Model Fidelity Trades

In order to demonstrate the impact of different modeling options on the output imagery, a suite of hyperspectral images cubes were created using different modeling options. The first major modeling difference focused on the inclusion of transition reflectance measurements in the material reflectance databases. This will be discussed in depth as part of the results discussion. The second modeling option was concerned with the amount of sub-pixel sampling that was utilized. A sub-pixel sampling rate of 1 x 1 results in pure pixels and sampling rates greater than 1 x 1 result in mixed pixels. The amount of oversampling impacts the clustering of the background classes and defines the minimum sub-pixel fraction that a target could manifest itself as in the image. For example, when using 3 x 3 sub-sampling the minimum target contribution is 1 out of 9 samples or approximately 11% pixel fill. The final parameter of interest was

the enabling of rigorous BRDF calculations. When the full BRDF calculation is enabled, the model attempts to quantify the incident irradiance from the hemisphere above the target rather than assume it is open, unobstructed, sky. When this calculation is performed, the model incorporates multiply bounced photons from background objects which results in phenomenological effects including “tree shine”. A summary of these simulations is captured in Table 1.

Name	Transitions	Sub-Pixel	BRDF
hydice0-1x	No	1x1	No
hydice1-1x	Yes	1x1	No
hydice2-2x	Yes	2x2	No
hydice2-3x	Yes	3x3	No
hydice3-2x	Yes	2x2	Yes

Table 1. Summary of simulated imagery that was created for algorithm evaluation.

3.2. Image Preparation

Before the image exploitation algorithms were applied to the data cubes, the images were subsampled to remove known atmospheric absorption bands. In addition, the first five (5) channels of the images had to be removed due to the limits of the field reflectance data acquired with the ASD field spectrometer. These noisy regions of the spectrum will be replaced by measurements performed with a stronger illuminate or with a lab instrument.

4. RESULTS

4.1. Spatial and Spectral Analysis

4.1.1. Spatial Analysis of Transition Regions

The results of the DIRSIG texture methodology can be seen in Figure 4 which contains a subsection of the synthetic image that contains grassy transition regions. Figure 4a shows an initial result where the supporting reflectance databases did not account for the variability present in the actual scene (see Figure 4c). Figure 5 illustrates where the field collection team measured a vary narrow reflectance region within the stressed and healthy grass. What was missing from these measurements was the reflectances in the overlap region where the reflectances of the healthy grass and stress grass were similar. Without these reflectances in the database, the transition between the two material types becomes very abrupt (see Figure 4a). To remedy this, we combined and synthesized the missing reflectances by mixing existing curves in each database. The resulting spectral database was used to produce the image in Figure 4b. Here, the curve selection algorithm is able to draw on a spectral database that contains all the relevant spectral nuances found in a grass region that transitions from a healthy to stressed state. Figure 4c is a high resolution aerial image of the same region used for comparison. This approach of generating hybrid databases was used for other materials as well.

4.1.2. Spectral Analysis of Transition Regions

In a hyperspectral image, the clusters associated with each material within the scene have potentially unique hyperdimensional shapes that cannot be represented with parametric models. Many hyperspectral clustering algorithms fail because they attempt to fit parametric models to data that is inherently non-parametric. One of the reasons that some algorithms perform better than expected on synthetic data is that the image generation models utilize similar parametric assumptions. As a result, the well-behaved modeled data can be exploited by these algorithms. One of the goals of this modeling effort was to demonstrate that synthetic image generation models can be built using non-parametric approaches to produce realistic data. The spectral clustering and correlation between one or more classes can be explored using ENVI’s⁸ n-Dimensional visualizer. This tool can project hyperspectral data onto a 2D plane and allows the user to interactively view the correlation and overlap between classes.

Two regions of grass (healthy and stressed) were analyzed in ENVI’s n-Dimensional visualizer. Figure 6 illustrates the projection of bands 31, 47, and 52 from the suite of simulations described in Table 1. The class clusters in Figure 6a further demonstrate the lack of overlap in the the hydice0-1x simulation discussed in the last section. The introduction of the missing transition curves results in the overlap observed in Figure 6b. However, since the model is producing pure



Figure 4. Simulated images of grass regions (soccer and track fields) showing (a) no transitions (b) with transitions and (c) high resolution aerial image for comparison.

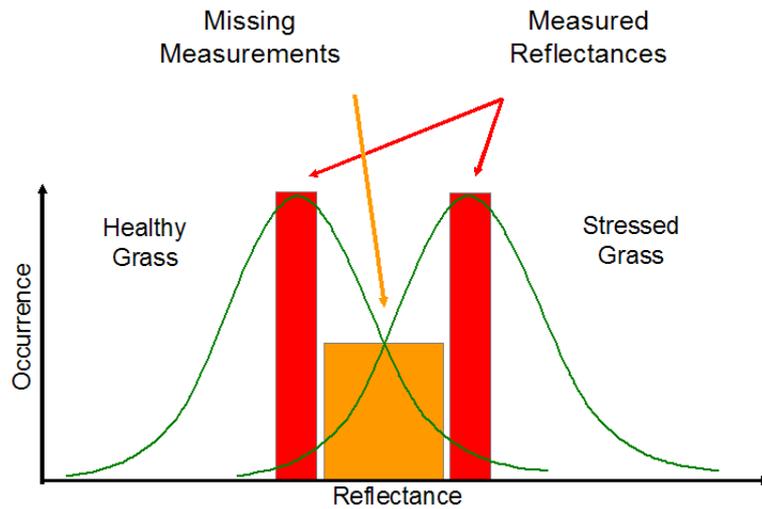


Figure 5. Hypothetical histogram showing distributions of healthy and stressed grass.

pixels created from a discrete set of reflectance curves for each material, the clusters appear as points in the hyperdimensional space. The introduction of pixel mixing in Figures 6c and 6d results in nearly continuous clusters. As the amount of sub-pixel sampling and spectral mixing increases the spectral clusters get narrower.

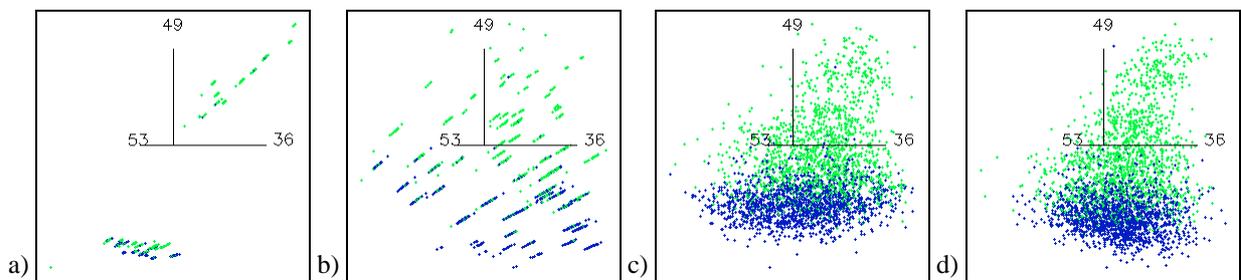


Figure 6. Projected grass class for the (a) hydice0-1x, (b) hydice1-1x, (c) hydice2-2x and (d) hydice2-3x simulations.

4.2. Algorithm Performance

The MegaScene and the produced imagery is not intended to be an exact match to a real data set but rather representative of a data collected from a real scene. To establish the robustness of the synthetic imagery, the performance of commonly used hyperspectral algorithms was evaluated in the context of commonly observed algorithm performance.

4.2.1. Classification

One of the first algorithm performance evaluations entailed the application of an unsupervised classification algorithm. We used the k-means¹¹ clustering algorithm because of its simplicity and widely understood nature. The algorithm was used to find seven (7) classes with two (2) to four (4) iterations to re-calculate class means (see Figure 7).

The algorithm was run on the hydice0-1x data set which contained no BRDF effects, no grass transition regions and pure pixels (see Figure 7(a-c)). We see from the results that the algorithm has adequately identified and clustered the major classes found in the original material map (see section 2.2). The excellent performance was expected because the spectral “clutter” in the image cube is very well behaved without the advanced modeling options enabled. The pure pixels in the imagery are easy to categorize into a class even with the simple minimum distance technique which is at the heart of the k-mean algorithm. Since Figures 7a and 7b are almost identical, the algorithm has converged after three iterations. The k-means algorithm was also applied to the hydice3-2x data set (see Figure 7(d-f)). This image cube was generated with the BRDF effects enabled, the enhanced grass transition databases and 2 x 2 pixel sub-sampling. The performance of the algorithm is significantly different from those observed with the simple simulation. Many regions, including various types of grass and asphalt, are misclassified due to the BRDF effects and the spectral mixing present in the data. Additionally the results show that the clustering algorithm found it more difficult to adequately categorize all the pixels after four iterations, while the previous data set converged after three.

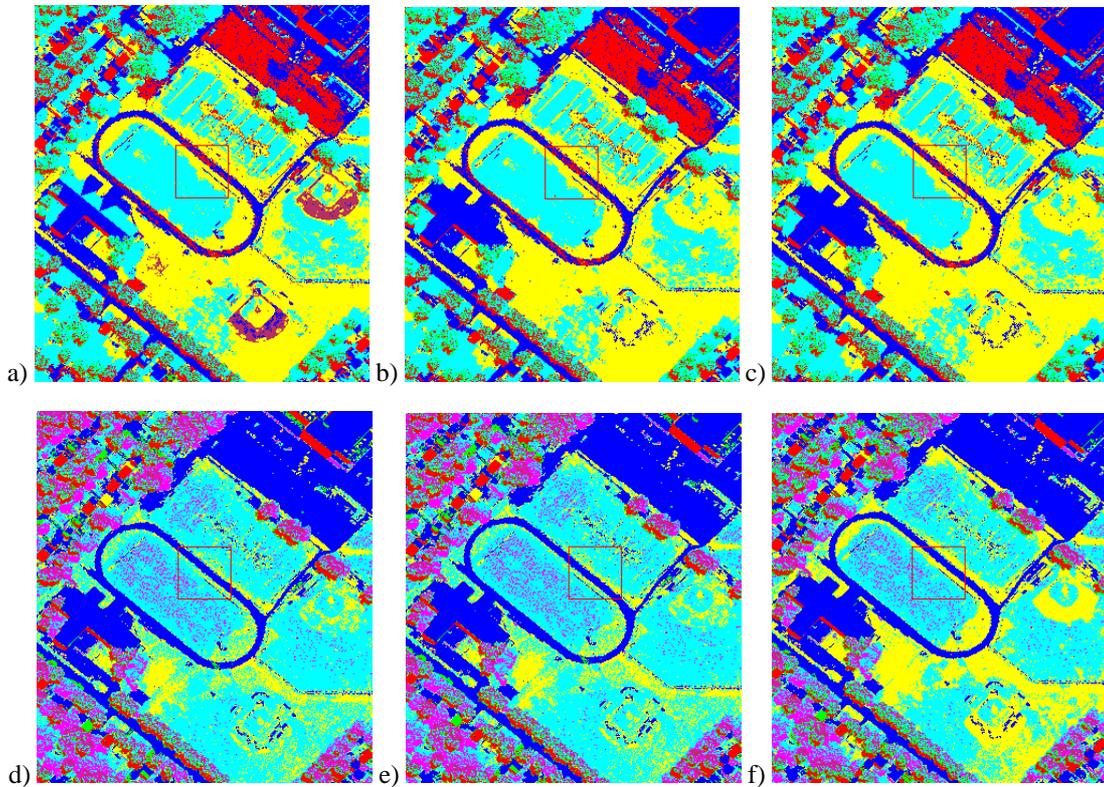


Figure 7. Results of k-means clustering algorithm. All images have 7 classes. Images (a-c) are results from the hydice0-1x data set with (a) 2, (b) 3, and (c) 4 iterations. Images (d-f) are results from the hydice3-2x data set with (a) 2, (b) 3, and (c) 4 iterations.

4.2.2. Target Detection

To explore the utility of the simulated imagery in the evaluation of target detection algorithms, we tested a pair of off-the-shelf algorithms contained in ENVI on panel targets that we deployed in the scene. The panels were assigned a man-made spectrum that has vegetation characteristics. Three panels measuring 1/2 meter, 1 meter and 2 meters in size were inserted into the scene resulting in a total of twenty-four (24) panels (eight (8) of each size). The panels were placed in both wide open and shaded locations (see Figure 8).



Figure 8. (a) Three panels in a parking lot area (b) a single panel in a tree shadow and (c) a single panel in a hard shadow behind a building.

The first algorithm that was tested was the Spectral Angle Mapper (SAM)⁹ algorithm. This algorithm classifies individual pixels in the scene by treating the hyperspectral signatures as hyperdimensional geometric vectors and performing a dot product of the pixel spectrum of interest to the target spectrum. If the resulting angle is below a user defined threshold, the pixel is classified as target. This algorithm is extremely fast but performs poorly when the pixel spectrum undergoes non-linear scalings due to illumination load changes (*e.g.* shadowing) or contamination due to spectral mixing. The performance of this algorithm is generally well understood which is why it was included in this demonstration. The second target detection algorithm was the Matched Filter (MF).⁹ The Matched Filter belongs to a general class of algorithms that attempt to find pixels similar to the target spectrum by projecting the image into a hyperspectral space that maximizes the separation between the target spectrum and the background.

The results for both target detection algorithms is summarized using receiver-operator characteristic (ROC) curves. The ROC curve describes the relationship between the probability of a true detection and the probability of a false alarm. The performance of a given algorithm will be affected by many parameters including the general spectral complexity of the scene, the similarity of the target to the background, the size of the target, etc. An ideal algorithm would approach a 100% detection rate while accruing zero false alarms (0%). In the absence of the ideal algorithm, the algorithm development community works to create algorithms that move the “shoulder” of these curves towards ideal performance (the upper-left side of an ROC plot). Historically, algorithms which are generally accepted as “poor” have been shown to perform ideally on synthetic data, which further demonstrates the lack of realism in the data. Within the confines of known and accepted physical processes, the modeling community works to create imagery that moves the ROC curve shoulder in the opposite direction from ideal detection (lower-right). This region is synonymous with the performance of algorithms that operate on real world imagery.

To create a ROC curve, the analysis tools must have access to truth data that defines where the targets are in the imagery. When real world data is utilized to measure algorithm performance, the truth maps are generated using field-gathered ground truth measurements which are then mapped to pixels within the scene. Due to the inherent uncertainty of the target location and the pixel/target overlap, the detection analysis can only be performed on a per target basis. For example, when a detection within an acceptable distance from the estimated location of the target occurs, the event is recorded as a positive detection. Additional detections of the same target are discounted.

The ROC curves were generated using a truth map that was derived from the truth maps created by the DIRSIG model during the simulation. This truth map allows us to create pixel-based ROC curves rather than target-based ROC curves. Pixel-based ROC curves can be created because the precise location and pixel overlap of the target is known. This pixel-based approach allows algorithm testers to explore minimum pixel fill detection rates. The target spectrum for the

panel was image derived from a pure pixel in the “hydice0-1x” simulation. The ROC curves in Figure 9 demonstrate the performance of the Spectral Angle Mapper and Matched Filter algorithms on the synthetic data sets. For each algorithm, a ROC curve was generated using different amounts of sub-pixel sampling. As the sampling rate increases, the target can potentially become a smaller fraction of the total pixel spectrum. The initial detection rates (low probability of false alarms) in this demonstration are expected to high because many of the target panels are fully resolved at the 1 meter ground sampling distance.

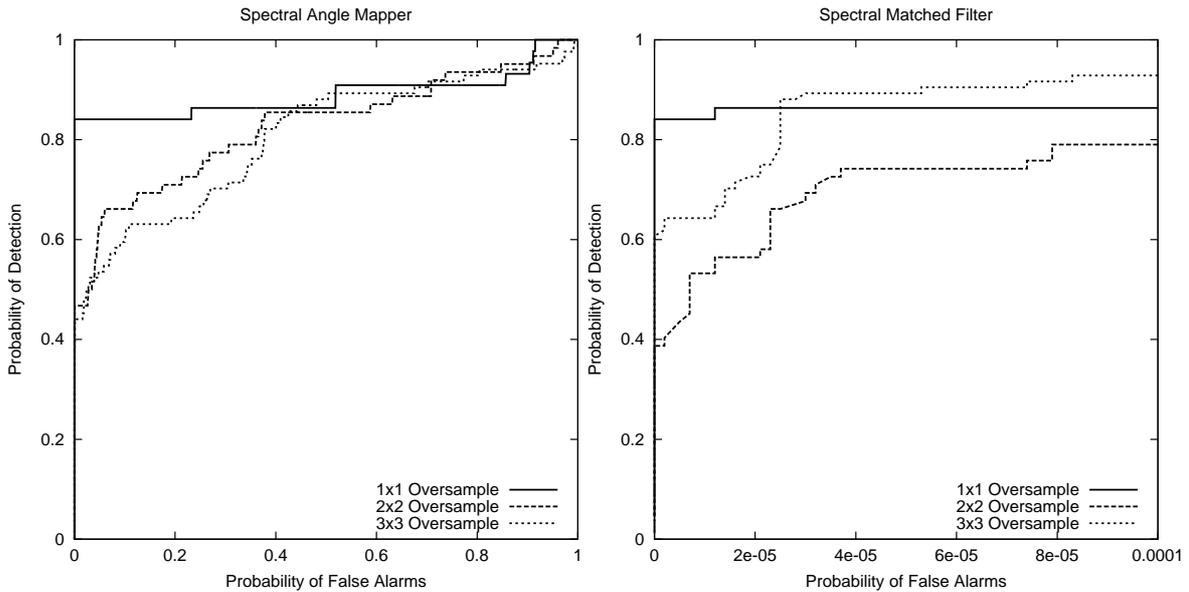


Figure 9. ROC curves for Spectral Angle Mapper (left) and Spectral Matched Filter (right) using different sub-pixel mixing amounts (note different Pfa axis scales).

The results for the SAM algorithm behave as expected. As the amount of sub-pixel sampling increases the ability to detect the targets decreases in an ordered fashion. The MF results are not initially intuitive since the performance of the algorithm does not always decrease as the sampling increases. Presently, we believe that the increased level of mixing “tightens” the hyperspectral clusters associated with the background classes which results in easier separation of the target from the background. The effect of mixing on the variance of the classes can be observed in the n-D visualizations in Figure 6. According to the Central Limit Theorem, the variance will eventually converge as the sub-pixel mixing increases. As a result, the performance trend is expected to behave similarly to those observed with the SAM algorithm when higher sub-pixel sampling rates are utilized.

5. SUMMARY

The large area scene database described in this paper is the product of a comprehensive geometric construction effort coupled with a carefully planned and executed field collection to support the optical property databases. The resulting database has been demonstrated to provide a level of geometric and spatial-spectral complexity over a large area that we believe is ground breaking.

The initial demonstrations of the scene and the generated imagery have focused on evaluating the fidelity of the generated image sets and on quantifying the importance of several modeling fidelity trades available within the DIRSIG model. The imagery generated of the MegaScene by the DIRSIG model has been shown to contain complex hyperdimensional class clusters with the shape and overlap that is observed in real hyperspectral imagery. The simulated images also demonstrated algorithm performance trends similar to those observed running the same algorithms on similar real-world data. The importance of well constructed supporting databases built from well collected background reflectances was also explored. Finally, the impact of pixel mixing on algorithm performance was also demonstrated.

The algorithm performance demonstrations were performed on noise-free images. The impact of uncorrelated and spectrally correlated noise sources within the instrument system are believed to have a significant impact on the performance of spectral algorithms. A back-end instrument model like the HySIM model developed by Veridian would further enhance the validity of the generated imagery and resulting algorithm performance.

The second tile of the MegaScene region is nearing completion. The combined area of the two tiles will encompass an area that is approximately 1 km x 2 km. The next phase of this effort focuses on extending the modeling coverage into the 3 - 14 micron region. To enable this region, the thermodynamic properties of the materials will need to be finalized and the field collection team will need to augment the reflectance databases with the thermal region spectral emissivities.

REFERENCES

1. Schott, J.R., Brown, S.D., Raqueño, R.V., Gross, H.N. and Robinson, G., "An Advanced Synthetic Image Generation Model and its Application to Multi/Hyperspectral Algorithm Development", *Canadian Journal of Remote Sensing*, Vol. 15, pp. 99-111, 1999.
2. Sheltler, B.V., Mergens, D., Chang, C., Mertz, F., Schott, J. R., Brown, S. D., *et al.*, "Comprehensive Hyperspectral Image Simulation I: Integrated Sensor Scene Modeling and Sensor Architecture", *Proceedings of SPIE AeroSense, Algorithms for Multispectral, Hyperspectral and Ultraspectral Imagery VI*, Orlando FL, Vol. 4049, pp. 94-104, April 2000.
3. Garland, M. and Heckbert, P. S., "Fast Polygonal Approximation of Terrains and Height Fields", Carnegie Mellon University Report *CMU-CS-95-181*, available online at <http://graphics.cs.uiuc.edu/~garland>, September 1995.
4. "Rhino: NURBS Modeling for Windows" software, information online at <http://www.rhino3d.com>
5. "Tree Professional" software, information online at <http://www.onyxtree.com>
6. Schott, J.R., Salvaggio, C., Brown, S. D., Rose, R.A., "Incorporation of texture in multispectral synthetic image generation tools", *Proc. SPIE Aerosense, Targets and Backgrounds: Characterization and Representation*, Vol. 2469, No. 23, 1995.
7. Brown, S.D., Raqueño, R., and Schott, J.R., "Incorporation of bi-directional reflectance characteristics into the Digital Imaging and Remote Sensing Image Generation model", *Proc. of the Eighth Annual Ground Target Modeling and Validation Conference*, pp. 163-171, August 1997.
8. "The Environment for Visualizing Images (ENVI)" software, information online at <http://www.ResearchSystems.com/envi>
9. Kruse, F. A., "The spectral image processing system (SIPS)— Interactive visualization and analysis of imaging spectrometer data," *Remote Sens. Environ.*, Vol. 44, pp. 145-163, 1993.
10. Stocker, A., Reed, I., Yu, X., "Multi-dimensional signal processing for electro-optical target detection, *Proceedings of SPIE*, Vol. 1305, pp. 218-231, 1990.
11. Tou, J. T. and Gonzalez, R. C., *Pattern Recognition Principles*, Addison-Wesley Publishing Company, Reading, Massachusetts, 1974.