THE SUPERCOMPUTING VISUALIZATION WORKBENCH FOR THE ANALYSIS AND CLASSIFICATION OF IMAGING SPECTROMETER DATA

Gregory S. Okin[1†]
Herb Siegel[2]
Jim Collier[2]
Craig D. Miller[2]
William J. Okin[3]
Dar A. Roberts[3]
Bruce Murray[1]
David W. Curkendall[2]
Thomas H. Painter[3]

1. INTRODUCTION

We have developed software for use with Airborne Visible and Infrared Imaging Spectrometer (AVIRIS) data in which user-oriented visualization tools and rigorous computational methods are interfaced for an intuitive analysis environment. The "Visualization Workbench" is a graphical user interface (GUI) that increases the ease and efficiency of sophisticated AVIRIS data analysis and provides a robust data exploration environment. This approach takes advantage of, but is not limited to, high-performance computation, visualization, and network technologies. It employs Multiple Endmember Spectral Mixture Analysis (MESMA) to classify AVIRIS data (Gardner, 1997; Painter et al., 1998; Roberts et al., 1998). Use of the Supercomputing Visualization Workbench effectively reduces weeks worth of AVIRIS data analysis to the course of a few hours or a day and enhances the user's ability to design scientifically-appropriate MESMA jobs through the integrated use of spectral attributes. It is a portable, modular tool that can grow with user needs, technique improvements, and computational advances. Please visit our website for a demonstration of the Visualization Workbench:

http://www.gps.caltech.edu/~arid/visualization_workbench/

In the past 30 years, digital spatial data— typically raster images— have become integral to a wide variety of the physical and natural sciences. Visualization of these data is the primary means by which we assimilate and communicate their information content and relies on the natural pattern recognition and training of scientists and analysts. Computation performed on these datasets allows deeper analysis and understanding but can often be time consuming and prohibitively complex. Combined computation and visualization tools that are sensitive to user needs is one way to enhance our ability to work with these datasets.

There is special need for the Visualization Workbench for work with very deep datasets like those from imaging spectrometers. Here, the number of bands vastly exceeds those needed for typical red-green-blue (RGB) display, the datasets themselves are cumbersome, computation can be very CPU-intensive, and analysis is often iterative or user-guided. Linked supercomputing and visualization are necessary for fast, concentrated analysis of imaging spectrometer datasets. The framework developed here may be applied to any intensive, interactive, or user-guided computational task that would greatly benefit from a concentrated, consistent, high-efficiency, and user-oriented analysis environment. Some examples of imaging spectrometer-related tasks that may be integrated with a Visualization Workbench are: reflectance inversion, coregistration and georeferencing, and radiative transfer or BRDF inverse modeling.

[1] Division of Geological and Planetary Sciences, 170-25 California Institute of Technology, Pasadena, CA, 91125.
[†] E-mail: okin@gps.caltech.edu
[2] Advanced Parallel Processing, Jet Propulsion Laboratory, 4800 Oak Grove Drive, M/S 126-234, Pasadena, CA, 91109-8099.
[3] Department of Geography, University of California, Santa Barbara, CA 93106.

The Supercomputing Visualization Workbench is a joint project between the California Institute of Technology (Caltech), the Jet Propulsion Laboratory (JPL), and the University of California, Santa Barbara (UCSB), and is made possible by the Caltech President's Fund.

2. GOALS AND CAPABILITIES

The primary goal of this project is to provide an analysis and data exploration environment in which an AVIRIS scene may be analyzed using MESMA to produce ground-truthable results in a few hours. The target user is computer-savvy but a non-expert in working with imaging spectrometer data. The approach is adaptable to other imaging spectrometer, multispectral, or multi-temporal remote sensing datasets.

The Visualization Workbench takes advantage of high-performance parallel computing and networking at JPL in order to create this analysis environment and to bring previously slow, CPU-intensive MESMA computation up to interactive speeds. This provides greater flexibility and user-guidance but does not necessarily limit the Visualization Workbench to use on a supercomputer embedded in a high-performance network. While this does make interactive speeds accessible, the user-oriented and intuitive aspects of the Visualization Workbench will still be available if it is running on a single-processor computer. For users with only remote access to a supercomputer, the visualization and GUI components can be piped through a 100 Mbit ethernet link to a modest desktop workstation (Silicon Graphics O_2) while computational tasks runs on the parallel machine.

The Visualization Workbench is not intended to be a stand-alone image processing package. Several very good such packages are already available commercially. The goal of this project is to enable a user to do a specialized task— MESMA analysis of AVIRIS data— very well and very fast.

3. OVERVIEW OF APPROACH

The Visualization Workbench consists of a set of portable stand-alone tools which may be divided into three major categories: graphical user interface, visualization tools, and computation tools. The tools communicate with each other by broadcasting information about user actions or system status to all other active tools. In this way, we have created a robust and expandable modular programming environment which can function as an incomplete set, if necessary, as well as one to which new features may also be added if desired.

The primary analysis technique employed by the Visualization Workbench is MESMA, an extension of spectral mixture analysis (SMA). SMA is based on the assumption that the reflectance spectrum derived from an air- or spaceborne sensor can be deconvolved into a linear mixture of the spectra of ground components, frequently called spectral endmembers. The best-fit weighting coefficients of each ground component spectrum, which must sum to one, are interpreted as the relative area occupied by each component in a pixel. MESMA is simply a SMA approach in which many possible mixture models are analyzed in order to produce the best fit (Gardner, 1997; Painter et al., 1998; Roberts et al., 1998).

In the MESMA approach, a "spectral library" is defined which contains spectra, convolved to the 224 AVIRIS bands, of plausible ground components. A set of m mixture models with n ($n \geq 2$) endmembers from the library is defined, with shade always present as one endmember in the model. The weighting coefficients (*i.e.*, fractions) for each model and each pixel are determined such that the linear combination of the endmember spectra produces the lowest RMS error when compared to the apparent surface reflectance for the pixel. Fractions are constrained to be between zero and one, and a valid fit is restricted to a maximum preset RMS error. Models that meet these constraints are recorded, which typically yields several possible models for each pixel.

Since the MESMA approach requires each pixel in an AVIRIS image be modeled by many different models, it can be very CPU-intensive. This fact makes it an ideal candidate for parallelization. In addition, MESMA processing is fundamentally user-guided since the models must be defined by the user to address a specific scientific question. Thus, the Visualization Workbench software aims at combining these two aspects of MESMA analysis into a single easy-to-use and interactive supercomputing interface.

The "RMS image" discussed below has dimensions 512 by 614 (the dimensions of an AVIRIS scene) by the number of models run (*m*). Each layer in the RMS image, therefore, gives the RMS error for a single model for each pixel in the AVIRIS scene and records those pixels not modeled within the given constraints.

As a final step, the one model for each pixel with the lowest RMS can be identified. This allows different pixels in the image to be modeled by different sets of endmembers. The optimal fractions of each endmember for this best-fit model are then recorded in the "fraction image". This image has dimensions 512 by 614 by ($n + 2$). The layers in the fraction image are: (1): the model number which gave the best fit for each pixel; (2): the RMS for that model; (3) through ($n+1$): the optimal fractions for each non-shade endmember, and; ($n+2$): the optimal fraction of photogrammetric shade.

This approach requires an extensive library of field, laboratory, or image spectra, where each plausible ground component is represented at least once. Including more than one spectrum of a ground component allows for the considerable spectral variability often found in desert vegetation, thus overcoming a difficulty identified by Franklin *et al*. (1993) of doing remote sensing in arid regions.

The spectral library is also associated with a "spectral attribute database" where relevant attributes of each spectrum in the library are maintained. These include the type of object that the spectrum is of (*e.g.*: shrub, grass, mineral), the genus, species, and phenological stage of vegetation spectra, characteristics of soil spectra, and other information about location, date of acquisition, and processing.

The spectral attribute database is vital to the Visualization Workbench approach as it enables the user to think about spectra in terms of the objects that the spectra are of instead of merely an indexed list of numbers.

4. HARDWARE COMPONENTS

The Caltech/JPL supercomputing assets and infrastructure enables interactive-speed computation capability in the Visualization Workbench.

The Hewlett-Packard Convex SPP-2000 System provides massively parallel processing capabilities to this research. This Examplar X-Class system is a multi-domain system consisting of 256 processors, each with a peak performance of 720 MFLOPS (10^6 floating point operations per second) for a peak system performance of 184 GFLOPS (10^9 floating point operations per second). The system has a total of 64 gigabytes of memory and 1 terabyte of disk. The network connection between Caltech and JPL is currently an OC-12 (620 Mbits/s) system and is being upgraded to a four-channel 100 Mbytes/s HiPPI network which will provide 400 Mbytes/s communication between JPL and Campus.

The Visualization Center at JPL has two Silicon Graphics (SGI) Dual Power Onyx graphics computers each of which has two Infinite Reality (IR) graphics engines. Each IR drives an Electrohome 8500 projector to provide a 60” by 80” display with up to 5 megapixels. The resulting four-screen display is housed in building 126 at JPL and is known as the "PowerWall". 576 gigabytes of (expandable) fiber channel RAID disk is attached to this system.

5. SOFTWARE COMPONENTS

Portability and modularity is a major goal for the Workbench. Portability is achieved by using ANSI standard C, Perl TK, Java, the Message Passing Interface (MPI), and other standard programming libraries. Use of the standard MPI libraries allows for easy rehosting of the system to other supercomputers or even to a single processor workstation. This makes the software available to a larger number of people on a variety of computers and also ensure that the software system can migrate to more advanced platforms as they become available. Modularity has numerous benefits among which are ease of modification and extensibility; isolation of the computational algorithms into a single library allows a researcher to create or modify them without knowledge of technical details concerning supercomputing, visualization, or network communications.

The Visualization Workbench assumes a hierarchical data structure. The primary directory contains an AVIRIS reflectance image, as well as a spectral library and its spectral attribute database. Invocation of the GUI while in this directory will allow the user do MESMA processing and visualization of the AVIRIS image and spectral library stored here. Within the primary directory are directories created by the Workbench— using names given by the user— in which MESMA input files, RMS images, and fraction images are stored.

5.1 Graphical User Interface

The GUI for the Visualization Workbench is the main object through with the user interacts with the Visualization Workbench. It has three main functions: View, Edit, and Submit.

The View function enables the user to view AVIRIS reflectance images, RMS images, or fraction images using the Digital Light Table (DLT) software (see below). Clicking on "View" in the main GUI, and then on one of these other options in the "View window" invokes a DLT window with the requested dataset. The View window also displays textual information about pixels that a user clicks on in a DLT window. This information includes: the image coordinates in the reflectance image; the RMS value, model number, and spectral attributes of the endmembers in that model in the RMS image; the model number, spectral attributes, and fractions in the fraction image.

The Edit function enables the user to name, create, and edit MESMA jobs by using the spectral attribute database. Clicking on "Edit" in the main GUI invokes the primary tools for creating MESMA jobs. This tool automatically queries the spectral attribute database, groups all spectra with identical attributes, organizes the groups hierarchically, and then presents the hierarchical list in three identical columns. Each group in each of the columns has a checkbox to its immediate left. Thus, all "Shrub" spectra occur together and within this, all "*Atriplex*" spectra and within this, all "*polycarpa*" in one group and all "*canescens*" in another. Checking any number of spectral groups in the first column only will define a two-endmember model job (shade is automatically included). Checking spectral groups in the first two columns will define a three-endmember model job, and checking spectral groups in all three columns will define a four-endmember model job. Clicking "Continue" shows a listing of all permutations (*i.e.* models) of the spectra in each spectral group which the user can edit at will. Clicking "OK" creates the MESMA input file from the modified list of models.

Since each group may represent many spectra, the combinatorial nature of the model creation process can easily create hundreds or thousands of models which must be run on each pixel. The user-guidable, attribute-based approach used in the Visualization Workbench helps the user make smart decisions about model creation. This allows the user to reduce the model set to a smaller, more appropriate one, thereby saving analysis and CPU time and improving the final results.

The Submit function enables the user to submit a MESMA job created by the Edit function to a remote supercomputer. Clicking on "Submit" in the main GUI window invokes a dialog box in which the user can define the user I.D., computer, and number of nodes to be used for the job. Clicking "OK" initiates a script that transfers all necessary information to the supercomputer, runs the MESMA job using the specified number of nodes, returns the results, and notifies the user that it has finished.

5.2 Visualization Tools

Currently, two different visualization tools work together in the Workbench to enable viewing and querying of data: the Digital Light Table (DLT) and a Gnuplot plotting tool.

Designed and implemented by the JPL Advanced Parallel Processing group, the DLT is a program used to interactively pan and zoom very large multispectral image files on SGI platforms with a single- or multiple-screen display. The DLT is used by the Workbench to display reflectance, RMS, and fraction images using user-defined band, model numbers, or fractions to create displays in RGB or grayscale. The large size of AVIRIS datasets, especially multitemporal ones, and RMS images necessitates innovative display capabilities in order to take full advantage of this high-dimensional data. DLT is perfectly suited to this task.

The Image Tile Pyramid is at the heart of DLT, and is how data for DLT is stored to disk. To create the tile pyramid for a particular data set, it is necessary to run the original data through the tiler program. This program takes the original data and breaks it up into tiles for DLT. This approach allows the DLT to perform well by caching only enough data to keep the output display updated. It only costs 1 1/3 the disk space required by the original input data. This allows disk files of hundreds of gigabytes to be viewed smoothly and interactively by DLT, and keeps memory usage down. Dataset tiling is an automatic feature of the Submit script.

Some mixture models tend to model the same areas equally well. The ability to see how mixture models cluster in the areas that they model is an extremely important piece of information in that it allows identification of redundancies and therefore more reliable image classification. As an analytical tool, DLT is extremely useful by allowing the scientist to view RMS images interactively in order to choose the best models and to establish model priorities.

The Gnuplot plotting tool is a modification of the readily-available Gnuplot graphing tool to plot spectral data. It is invoked by clicking on a pixel in a DLT window. Clicking on a pixel in the reflectance image produces a plot of the reflectance spectrum of the chosen pixel. Clicking on a pixel in the RMS or fraction image produces a plot of the reflectance spectrum of the chosen pixel, the library spectra of the endmembers in the active model[*], and the residual spectrum (reflectance spectrum minus best-fit model spectrum) for that model and pixel. In this way the user has a full suite of spectral information readily available. This information can be used to modify the model lists for subsequent MESMA runs in order to improve final results.

5.3 Computation Tools

The primary computation tools used in the Visualization workbench is Multiple-Endmember Spectral Mixture Analysis (MESMA). MESMA is comprised of two computational algorithms.

The first MESMA algorithm is used to produce RMS images from reflectance images. Since every model must be computed for every pixel independently, this algorithm is intrinsically parallelizable. Our implementation of this algorithms runs on any supercomputer running MPI and displays acceleration proportional to the number of nodes used. Running it with 256 nodes, therefore, is roughly two orders of magnitude faster than on a single-processor machine.

The second MESMA algorithm is used to create a single fraction image from the RMS image by sorting though the models for each pixel, choosing the one model with the lowest RMS, and recomputing the endmember fractions that give the best fit with the reflectance spectrum of that pixel. This algorithm is run once on each pixel and therefore is relatively fast. It has not, therefore, been parallelized and runs on the local single-processor machine.

The MESMA approach utilized by the Visualization Workbench is still being improved. Due to its modular nature, new improvements to the MESMA approach can easily be assimilated into the Workbench simply by making small changes in the necessary scripts that point to new code. The basic image framework—reflectance, RMS, and fraction images— is not likely to change.

6. SUMMARY

Processing and visualization of imaging spectroscopy is typically computationally intensive. We have presented a report on our efforts to use the Caltech's High Performance Information Technology infrastructure to

---

[*] The active model for an RMS image is that model number displayed in the red channel by the DLT. For a fraction image, it is the model that best fit the spectrum for that pixel.

place AVIRIS data processing within the context of a unified supercomputing, high-bandwidth network, and high-performance visualization structure. Any computationally-intensive data processing requiring user guidance may benefit from this general framework.

MESMA, similar to traditional spectral mixture analysis but differing in that many different permutations of spectral endmembers are applied to each pixel in a scene, is our principle computational method. Because several hundred models are often run on every AVIRIS pixel, this technique is computationally intensive and requires user guidance and oversight. We have developed a version of MESMA that runs on a distributed computer linked with a unique visualization system.

Thus, the Visualization Workbench provides a desirable and timely tool to link massive computation capability with high-performance visualization tools. In this framework, computation may be executed and displayed at interactive speeds to allow concentrated and detailed user-oriented analysis. The GUI allows display of AVIRIS images, design and submission of MESMA jobs to a remote supercomputer, and display of results in a manner that is sensitive to user needs. The resulting analysis and data exploration environment is highly interactive and allows the user to effectively guide the multiple iterations required in this type of analysis.

Use of the Supercomputing Visualization Workbench effectively reduces weeks worth of AVIRIS data analysis to the course of a few hours or a day and enhances the user's ability to design scientifically-appropriate MESMA jobs through the integrated use of spectral attributes. It is a portable, modular tool that can grow with user needs, technique improvements, and computational advances.

7. CITED REFERENCES

Franklin, J., Duncan, J., and Turner, D.L., 1993, Reflectance of vegetation and soil in Chihuahuan desert plant communities from ground radiometry using SPOT wavebands: Remote Sensing of the Environment, v. 46, p. 291-304.

Gardner, M., 1997, Mapping Chaparral with AVIRIS Using Advanced Remote Sensing Techniques [Master of Arts thesis]: Santa Barbara, CA, University of California, CA.

Painter, T.H., Roberts, D.A., Green, R.O., and Dozier, J., 1998, The effect of grain size on spectral mixture analysis of show-covered area from AVIIRS data: Remote Sensing of Environment, v. 65, p. 320-332.

Roberts, D.A., Gardner, M., Church, R., Ustin, S., Scheer, G., and Green, R.O., 1998, Mapping chaparral in the Santa Monica Mountains using multiple endmember spectral mixture models: Remote Sensing of Environment, v. 65, p. 267-279.