

Scientific Visualization for the Mars Exploration Rovers

Mark W. Powell, Jeffrey S. Norris, Marsette A. Vona, III,
Paul G. Backes and Justin V. Wick
*Jet Propulsion Laboratory
California Institute of Technology*

Abstract—We present an account of the scientific visualization techniques that were used to support science planning for NASA's 2003 Mars Exploration Rover Mission. The great sophistication of the rover's Athena Science Payload required a wide variety of visualization modalities to make the large amounts of scientific data accessible and readily understood by the science planners under the tight time constraints of the tactical planning schedule. A number of techniques were also newly developed to support science operations, including on-the-fly mosaic stitching and rendering, image cube visualization, and data fusion. Many other well-established visualization techniques were applied to science operations also, such as anaglyph stereo, high performance 3D visualization, and optimized I/O and memory management techniques for the loading, processing, and visualization of very large data sets. We discuss each of these topics and demonstrate their application in the MER mission with examples.

Index Terms—Visualization, telerobotics, science planning, image processing, coregistration

I. INTRODUCTION

The NASA's 2003 Mars Exploration Rover Mission brought arguably the most sophisticated set of scientific instrumentation to the surface of Mars in the history of exploration on that planet: the Athena Science Payload. These instruments are collectively capable of observing the Martian environment in stereo, color and infrared, quantifying the elemental chemistry of rock-forming minerals, characterizing the oxidation state of ferrous rocks and soil, and grinding beneath the weathered exterior surfaces of rocks to reveal their history by examining their interior structure. Conducting the operations of this historic mission of exploration required a collection of the most effective and efficient scientific visualization techniques available, as well as the development of several new techniques to encompass the full range of capability of these science instruments. In this work we present an account of the visualization techniques used in science operations and the rationale for their use. We also present examples of how these techniques were applied and our conclusions as to what extent they were valuable.

A. Downlink Data Products

Communication from the rovers on Mars to Earth occurs in discrete windows of opportunity. The number and length of these windows are defined by combinations of: time period of line of sight availability between the rover and Earth or one of the orbiting communication relays aboard Mars Global Surveyor or Mars Odyssey,

the relative position and orientation of the rover on the surface relative to Earth or a given orbiter, the availability of solar array and battery power, and thermal limitations of the rover's transmitters and other subsystems. The cumulative available bandwidth for a typical sol (Martian day) of mission operations is about 30 megabytes from each rover. In order to obtain as many images and spectral observations as possible, image subframing (cropping to less than a full image frame), downsampling, pixel scaling and compression (either lossless or lossy depending on the acceptable accuracy of the observation) may be used. For example, in the acquisition of three-color imagery of a particular area, it is often convenient to capture one of the colors at something close to full spatial resolution (1 megapixel, 12-bit samples), and the other two colors at a downsampled and reduced pixel scale, as the quality of the color information is not appreciably diminished in many cases. Similarly, multispectral data products are sometimes compressed depending on the acceptable fidelity needed for a given observation.

B. Spacecraft Activity Planning

Most activity planning for the rovers is tactical, i.e. based on the most recent images and other science data received from the rover. For instance, if a robotic arm placement was to have been completed and the arm stowed afterward, and we next intend to drive to a new location on the next sol, confirmation of the arm stow via telemetry is required in order to eliminate potential risk to spacecraft health. After the drive, accurately pointing the remote science instruments at targets requires the use of recently downlinked imagery of the surrounding environment for context, as the rover's localization knowledge may be inaccurate after a drive if localization refinement algorithms such as visual odometry or sunfinding have not been executed on board.

Tactical activity planning is conducted on a per-sol schedule, where all the activities that the rover will carry out on the next sol are created on the current sol, refined, scheduled, simulated, and safety checked prior to uplink to the spacecraft on the morning of the next sol. All of the steps that must occur after the creation of activities take approximately 7 hours with a 2 hour margin before the spacecraft expects to get the command sequence for that sol (at which point if it is not received, the entire sol would likely be wasted). This leaves approximately 4-6 hours for the entire science team to review the latest data, produce scientific analyses to conclude the success

or failure of the most recently conducted science activities, and create a high-level science activity plan for the coming sol.

To serve the needs of this demanding timeline, high-performance data visualization is a key requirement for the effective planning of science observations. We begin with an overview of the rover's instrument suite. In the following sections we will discuss the various scientific visualization techniques that were used. We present our results and conclude with a discussion of the relative usefulness in practice of each of these visualization techniques.

II. MER INSTRUMENTATION

The Athena Science Payload [1] consists of two instruments for conducting remote science observations and four instruments for in situ science observations. The remote science instruments are the:

- Panoramic Cameras (Pancam): a stereo camera pair with a field of view of 16.8 degrees mounted on a pan/tilt mast (the Pancam Mast Assembly, or PMA). Both cameras have an actuatable filter wheel, making this camera pair capable of acquiring band pass images in 11 unique color spectral bandpasses in the range of 400 to 1100 nm.
- mini-Thermal Emission Spectrometer (Mini-TES): an infrared spectrometer, also targetable through the PMA, that acquires remote spectra in the range of 5 to 29 μm .

For observations in close proximity to or in contact with the Martian surface, there are four instruments on each rover mounted on a robotic arm called the Instrument Deployment Device (IDD). The in situ instruments include the:

- Microscopic Imager (MI): a dichromatic camera that is deployable to capture high resolution (30 μm per pixel) images of rocks and soil.
- Alpha Particle X-ray Spectrometer (APXS): an alpha and X-ray spectrometer that can detect the abundance of all elements involved in the formation of rocks except hydrogen.
- Mössbauer Spectrometer: a gamma spectrometer that is suited to detect the oxidation state of iron, such as when it is altered by the presence of water.
- Rock Abrasion Tool (RAT): a rock grinding tool that removes material in a 4.5 cm diameter circular area to a depth of up to 5 mm, exposing fresh, unweathered rock surfaces for observation. The RAT also has a brush component that can gently remove surface dust from the outer surfaces of rocks.

The rover's engineering instruments include the:

- Hazard Avoidance Cameras (Hazcams): two monochromatic stereo camera pairs, one facing front and one facing to the rear, each with a 120 degree field of view.
- Navigation Cameras (Navcams): a monochromatic stereo camera pair mounted on the PMA to acquire stereo mosaics of some or all of the local terrain

immediately surrounding the rover with images of 45 degrees field of view each.

The images from all of these cameras are used in detailed traverse planning and sequencing to manually identify waypoints of interest and unsafe terrain (obstacles) to avoid. In addition, the front Hazcam imagery is used to map the work volume for the placement of the robotic arm for in situ observations.

III. STEREO IMAGE VISUALIZATION

Nearly all images returned from the Mars Exploration Rovers are stereo image pairs, as the rovers were designed to emulate a human geologist on Mars. To the geologist, particular morphologies and differences of textures such as layering in the environment speak to the history of the formation of the region. Stereo image pairs create for human analysts and algorithms alike a map of the morphology of any given area in view. The only cameras on the spacecraft that aren't stereo are the Descent imager, which is mounted on the lander and was used only to capture images of the local area from high above as the rover came to land, and the MI. The capability of the IDD to deploy and position the MI at virtually any position and orientation in the work volume allows the science planner to position this camera to acquire two images of the same specific area from slightly different, adjacent points of view, and in this way by knowing the baseline distance and relative orientation of the IDD stereo range can be computed. Further, since the MI has a very short depth of field, and surface morphologies of scientific interest are often quite variegated, multiple images along the same direction of view at slightly varying distance are often taken to ensure that in at least one image any one part of the surface in view will be in optimal focus. As a result, depth from focus techniques can often be used to derive range maps of surfaces where these image "stacks" were acquired.

For manual analysis of terrain morphology, anaglyph stereo visualization provides a very effective way to present a stereo pair of images that allows the unconscious ability of the brain to derive depth from the pair of images where the baseline separation of the cameras that acquired the pair is similar to that of human eyes. By simply overlaying the left image of the stereo pair in red onto the right image in blue or cyan, an analyst wearing red-blue filtered glasses can intuit the shape of rocks and varying depth of the surface with little effort. Similarly, this effect can also be achieved by presenting a stereo image pair using stereo display capabilities supported in many modern graphics controllers coupled with shutter glasses that are synchronized with the images on screen, providing the left image only to the left eye and the right image only to the right. This capability is not universally available on all computer hardware, however, and also requires a display having a refresh rate of 120 Hz or better, in order that each eye may see stereo images, animations, or other 3D renderings at 60 Hz in each eye that appear smooth and realistic. Not

every CRT monitor is capable of providing such a high refresh rate, and LCD monitors are thus far incapable of providing this type of display. Thus, for laptop applications, many desktop hardware applications, and printed materials, anaglyph stereo is the preferred option. Fig. 1 shows an example of a red-cyan anaglyph of Columbia memorial station upon Opportunity's egress at Meridiani Planum.




Fig. 1. Rear Hazcam view of Columbia Memorial Station upon Opportunity's egress at Meridiani Planum in red-cyan anaglyph.

Every camera on the rovers is modeled both by its intrinsic parameters such the optical characteristics of its lens and field of view, and its extrinsic parameters, such as its relative position and orientation with respect to the rover body. The types of models used for the cameras are of type CAHVOR [2], which model the position, orientation, field of view, orientation of the imaging element array with respect to the optical axis, and radial lens distortion. For the particular case of the Hazcams which employ fish-eye lenses to capture a 120 degree field of view, the extended model CAHVORE [3] was used to model the parameters of fish-eye image distortion as well. Both of these models provide a mapping of 3D points in space to 2D image coordinates that are useful for visualization.

Science activities directed at points on the surface are whenever possible parameterized by very accurate 3D positions referred to as targets. These targets are based on correlating a stereo image pair into a range image or range map that records the distance along a ray emanating from every pixel of an image (by convention, the left image of a pair) to the surface in the scene. These targets are then named for their geological characteristics such as morphology, etc. to make them more familiar to scientists and to be useful feature names in science planning discussions. Visualization that serves such discussions can

display an image of the area containing a target with an iconic annotation of the target's position and a label displaying the target's name. The camera model of the image presented to the user provides the mapping from the 3D location to the appropriate pixel location. These targets can therefore be displayed as annotations not only on an image of the pair from which the target was first derived, but also images taken by other instruments of the same region, or even images taken of the same region at a later time after the rover has driven to a new location but is still able to view the region in question, provided one also takes into account the coordinate transform representing the difference between the rover's position and attitude at the time which the image pair was acquired in which the target was defined, and that of the rover at its new location. In general, any three dimensional shape representation can be annotated on a camera-modeled image by taking advantage of its inherent 3D to 2D mapping capability.

IV. MOSAIC VISUALIZATION

Mosaicking is the process of stitching together a collection of images of individual regions to form a more comprehensive rendering of the local environment. Mosaic visualization is useful for various tasks in science planning, such as prioritization of remote science targets or debating a number of available traverse paths in terms of safety and science potential. Planning discussions such as these require comprehensive mosaicked maps of the entire local environment in order to be conducted effectively. As we have discussed, the images used for tactical science planning are not downlinked from the rover until just before the beginning of the science planning process, so mosaic visualization must be fast (no more than several minutes) in order to keep pace, or else delay the tactical timeline. For this purpose, a real-time mosaic visualization technique was developed.

In order to render image mosaics on the fly (at whatever time new images arrive that are useful for tactical planning), a warping technique was developed that has real-time performance while still producing high-quality mosaics. Consider a set of images taken from a camera on the PMA where for each image the camera is slewed in azimuth (pan) and/or elevation (tilt) to point in a different direction. Imagine a sphere centered on the origin of the rotation frame of the pan-tilt mechanism. As each image has a corresponding camera model, and as each camera model maps 2D pixel locations to 3D rays emanating into the world. Each of these rays intersects the sphere at a particular point. If for every pixel in each image the brightness value is mapped onto the interior of our sphere, we obtain an image in spherical coordinates that effectively stitches together all of the images. We can then reproject the image from spherical coordinates to visualize it more effectively, such as a cylindrical projection. This the most often used map projection to present this type of imagery, as it realistically depicts what one would see if they stood in place of the rover on Mars and turned their head about. Fig. 2 shows an example of a cylindrical projection. Overhead

image projections are also sometimes useful in tactical science planning, and the spherically-projected image can also be reprojected (such as with a polar-azimuthal plot) to visualize the local environment from overhead. Fig. 3 shows an example of a polar-azimuthal mosaic. This mosaicking technique is further detailed in [4].




Fig. 2. Navcam cylindrical mosaic view from the Opportunity rover looking back after egress from Eagle Crater at Meridiani Planum.




Fig. 3. Navcam polar-azimuthal mosaic view of Eagle Crater from the Opportunity rover taken prior to its egress from the lander. The image is annotated with various target locations denoting features of interest and locations where the rover traversed to following egress.

V. 3D TERRAIN VISUALIZATION

Visualizing the local terrain of the rover in 3D is useful for many types of activities, including instrument arm deployment, navigation, remote imaging, and trenching. For these types of activities, 3D datasets derived from the rover's stereo cameras provide the highest fidelity coverage of the nearby surfaces. Operationally, placement of the

robotic arm instruments on or near the terrain requires sub-millimeter localization accuracy, and for this reason stereo range derived from front Hazcam stereo pairs is strongly preferred. For remote imaging or navigation, stereo range derived from Navcam or Pancam (if available) image pairs is useful if the rover hasn't moved significantly since the images were acquired, i.e. if the localization knowledge of the rover has not appreciably degraded. If the localization knowledge of the rover is poor, engineers can make improved estimates of localization by manual comparison of images taken before and after movement and update the rover's localization knowledge from the ground prior to further pointing of remote imaging instruments.

A. Visible Scalable Terrain (ViSta)

A representation of surface geometry, called Visible Scalable Terrain (ViSta), was designed to meet the needs of 3D terrain visualization derived from stereo range data [5]. First, range maps are computed by correlating a stereo pair of images. The range maps are then triangulated, producing a triangular mesh that represents the terrain surface to the accuracy limits of the correlator. Since, by convention, each range map point is computed along a ray from a pixel in the left image of the stereo pair, the left image can be used as a texture map for the mesh, reproducing the brightness (in the case of Pancam terrains, color) of the surface at the time that the images were acquired. The ViSta format includes a representation of the mesh geometry as well as the necessary texture mapping data (texture coordinates). Fig. 4 shows an example of a 3D view of a set 10 of ViSta terrains derived from color Pancam stereo imagery.




Fig. 4. 3D view of a set of 10 ViSta terrain segments derived from color Pancam stereo imagery of the rock outcrop known as El Capitan in Eagle Crater, Meridiani Planum.

B. Level of detail (LOD) management

In the case of MER camera instruments, each full-frame image is 1 megapixel, which in the case of a full correlated range map results in a geometry of 2 million triangles per image. To achieve full Navcam coverage of

the local environment, as many as 20 stereo image pairs are required, resulting in a set of geometry of nearly 40 million triangles in size. To that we may wish to add front and rear Hazcam coverage, and any number of Pancam color stereo terrain maps, involving up to 100 million triangles or more. In order to render these complicated datasets effectively on low-cost hardware, the ViSTa terrain representation includes levels of detail (LOD). Graphically, terrain maps of reduced LOD can be substituted for full-fidelity terrain when viewing the surface from a distance. This results in a much lower number of triangles needed to fully render the environment when the viewer is at a distance, viewing most or all of the local terrain. As the viewer approaches a particular segment of the terrain, the viewpoint frustum is used to selectively render only the terrain segment currently visible at an increasingly higher fidelity, potentially reaching the highest fidelity available when the viewer is nearest to it. This strategy helps ensure that the total number of triangles needed to effectively render the environment is on the order of 2-3 million instead of tens or hundreds of millions.

C. Target selection

As ViSTa terrains are intended to be used in science planning, holes in the range map (where the correlator fails to achieve a match between left and right images) are not filled, but left empty. To do otherwise would allow targeting science activities at 3D points on the surface that may be significantly inaccurate.

Our need to visualize the terrain using lower LODs when surfaces are far from the viewer necessarily involves geometric approximations to the highest fidelity range data available. In other words, if a target was picked by casting a ray from the viewer into the virtual scene and intersecting with the nearest triangle in the geometry, that intersection will not in general be an actual correlated point from the range map, due both to the LOD simplification and to trilinear interpolation across the intersected triangle. To eliminate such inaccuracy, the range map is always used as the authoritative data product for targeting. When a target is picked on a terrain, regardless of the level of detail currently being viewed, the (trilinearly interpolated) texture map coordinates of the intersection point are used to look up the associated 3D point from the range map at full accuracy. Using this general strategy to ensure the highest accuracy of the target, we have optimized our ray intersection algorithm. We begin not by intersecting whatever LOD is being displayed, but by first intersecting with the bounding boxes of each terrain in the view. For each hit, we intersect again with the corresponding lowest LOD triangle mesh, taking the nearest such intersection point to the viewer as the picked point. Finally, we look up the associated range point using the texture map correspondence, and we use that as the target. This algorithm covered in further detail in [6].

The data returned from the Mini-TES quite unlike the conventional images such as those from Navcam and Pancam in several important ways beyond the fact that it records only emissions in the infrared spectrum. In contrast to a CCD camera, where a fraction of second of exposure time records the entire observation of an area, the Mini-TES requires this amount of time (if not more) per pixel, and for each pixel the Mini-TES capture mirror must slew to a slightly different elevation and the PMA rotate as needed to direct the optical path of this instrument. The data itself is more difficult to interpret since not only are pixels quite large compared to those of the cameras, resulting in much lower spatial resolution, but each pixel records an infrared spectrum at 167 distinct wavelengths. Thus, the scientist's task is to parse through an image consisting of 167 spectral bands at low spatial resolution relative to the available contextual imagery of the region where the spectroscopy was conducted. Two techniques were developed to assist in the analysis of this type of data, image cube visualization and data fusion.

A. Image Cubes

In order to quickly identify individual spectral bands of scientific interest, the image cube is an effective visualization technique. The user is given a slider or similar GUI control that can move a plane through a virtual cube of the multispectral data along any of three axes: azimuth, elevation, or wavelength. In this way, the user can quickly dismiss spectral regions of noise or inactivity and readily identify useful wavelengths that deserve further analysis. The implementation of this visualization technique for the case of the Mini-TES was achieved by rendering a plane that is oriented orthogonally with respect to one of the major axes of the cube and texture map one spectral band of the data on it. The extent of any single Mini-TES in either azimuth or elevation is more than 100 pixels, so the texture map to construct is relatively small and can be converted quickly from the spectral data and rescaled to 8 bits per channel for rendering. It is also very valuable to a spectroscopy specialist to be able to plot line graphs of spectra by wavelength, as this is the way in which this data is conventionally presented and readily interpreted by an experienced analyst. Fig. 5 (upper left) shows an example of an image cube visualization of Mini-TES spectral data.

B. Data Fusion

In order to mitigate the issue of low spatial resolution of the infrared spectrometer data, it was necessary to develop a technique to overlay it onto contextual imagery of high spatial resolution such as that of the Hazcam, Navcam, or Pancam. We will now discuss the particular technique, referred to as data fusion. Just as it is straightforward to relate 3D points to image coordinates in camera-modeled images, so too can 3D points be mapped to a view frustum of the Mini-TES instrument, which is well-modeled by a cone extending outward from the aperture where incoming

light falls onto the capture mirror. Using knowledge of the rover's position and attitude at the time when the spectra were acquired, it is straightforward to determine whether a given 3D point lay within one or more view cones of the spectrometer. If more than one view cone contains the point, this may be resolved in one of several ways, such as taking the view cone where the point lay nearest to the axis of the cone, or to consider the weighted contribution of each of the set of cones to the observation of this surface point in proportion to its distance from each cone's axis. Once the cone or cones containing a particular point on the surface have been identified, the infrared data may be mapped onto one or the other image of a stereo image pair, using its range map to provide the 3D point needs to complete the mapping from spectrometer view cone to 3D point to camera pixel. Iterating this technique over every pixel in the range map produces a co-registered image of infrared spectral data that is easily overlaid onto the stereo image using alpha transparency to vary the opacity of either of the visible or infrared images as needed. Fig. 5 (upper right) shows an example of data fusion. This technique is covered in detail in [4], [7].

VII. PERFORMANCE

Although the bandwidth of data transfer from each of the spacecraft to Earth is on the order of 30 megabytes (depending on the availability and configuration of orbital communication assets), the amount of derived data products based on the telemetry is often 4-5 gigabytes. In order to effectively visualize this large amount of data within the constraints of the science planning timeline, it is necessary to maximize performance. We will now discuss our performance strategy in the areas of I/O and memory management.

A. I/O

In general, computer applications that feel responsive to the user are those that provide responsive feedback for any given operation with minimal delay. To provide responsive feedback to users loading large datasets, it is necessary to perform I/O asynchronously and in a multithreaded fashion. Asynchronous I/O allows the visualization to incrementally render segments of the data product as they are loaded, until they are all completely loaded into memory. The value of this feedback to the end user cannot be understressed: a modern computer application will be judged to be either fast or slow by a user depending on whether they receive near-immediate feedback.

A multithreaded strategy for I/O is advantageous in several respects. Executing lengthy operations in a GUI event thread creates unreasonable feedback delays. Performing I/O in a thread or threads dedicated to that task allow other threads that need not block on I/O execute, enhancing overall application performance.

For large sequentially-organized files such as images, terrain maps and flat-file databases, memory-mapped files

an invaluable tool for improving I/O performance. This technique allows programs to buffer files for reading and writing into main memory using the underlying O/S virtual memory paging system (which is already highly optimized) to read or write pages of the file to or from disk as required. This can result in substantial performance benefits when files are organized and accessed in a sequential fashion. Querying range maps for regions of interest from a corresponding camera image loads just the pages containing range values in the desired region, without loading the entire file. Loading of ViSta format terrain files also benefits substantially from this strategy, as the levels of detail of the geometry are organized in a sequential fashion within the file; higher levels of detail are never loaded into memory if they are never needed. Creation of data products such as these can be similarly optimized by this technique, writing a large dataset to a memory buffer and then finally paging to disk with the virtual memory manager.

B. Memory Management

Very large images, such as mosaics of size 10,000 pixels wide or more, often cannot be represented in their entirety within memory constraints. To manage such large datasets in memory, a tile-based image cache in memory is required. A cache of tiles (sub-images) is maintained in memory and populated on-demand as user requests are made to load, process, and render images. In this fashion, it is possible to either zoom out to view a very large image in its entirety (albeit at a reduced scale), or zoom in to view the image at its full resolution and pan around to inspect one part of it at a time. Tiles are computed on-demand and inserted into or removed from the memory cache as needed to service the current needs of the application. For image processing, operators must be specially written to compute results on a tile-by-tile basis. In order to provide 3D immersive views while conserving as much memory as possible, texture maps are generated first from any images that may already have been loaded for 2D visualization before attempting to load them from storage, and in this way redundant in-memory copies of the same image data are avoided.

VIII. CONCLUSIONS

All of the visualization techniques were made available to the Athena Science Team as part of the Science Activity Planner (SAP) software application. The science team participants used SAP as their primary tool for downlink data analysis. The overall high performance of all of the aforementioned visualization techniques was very much welcomed by the scientists. Indeed, had any of these visualization tools required excessive time, they would not have been useful in tactical science analysis.

The Athena Science Team represented a broad cross-section of the science community, consisting of geologists, geochemists, spectroscopy and mineralogy specialists, and atmospheric scientists who formed the major groups involved in tactical science planning. The science team as a whole favored those visualization techniques that presented information in 2D most of all. Visualization

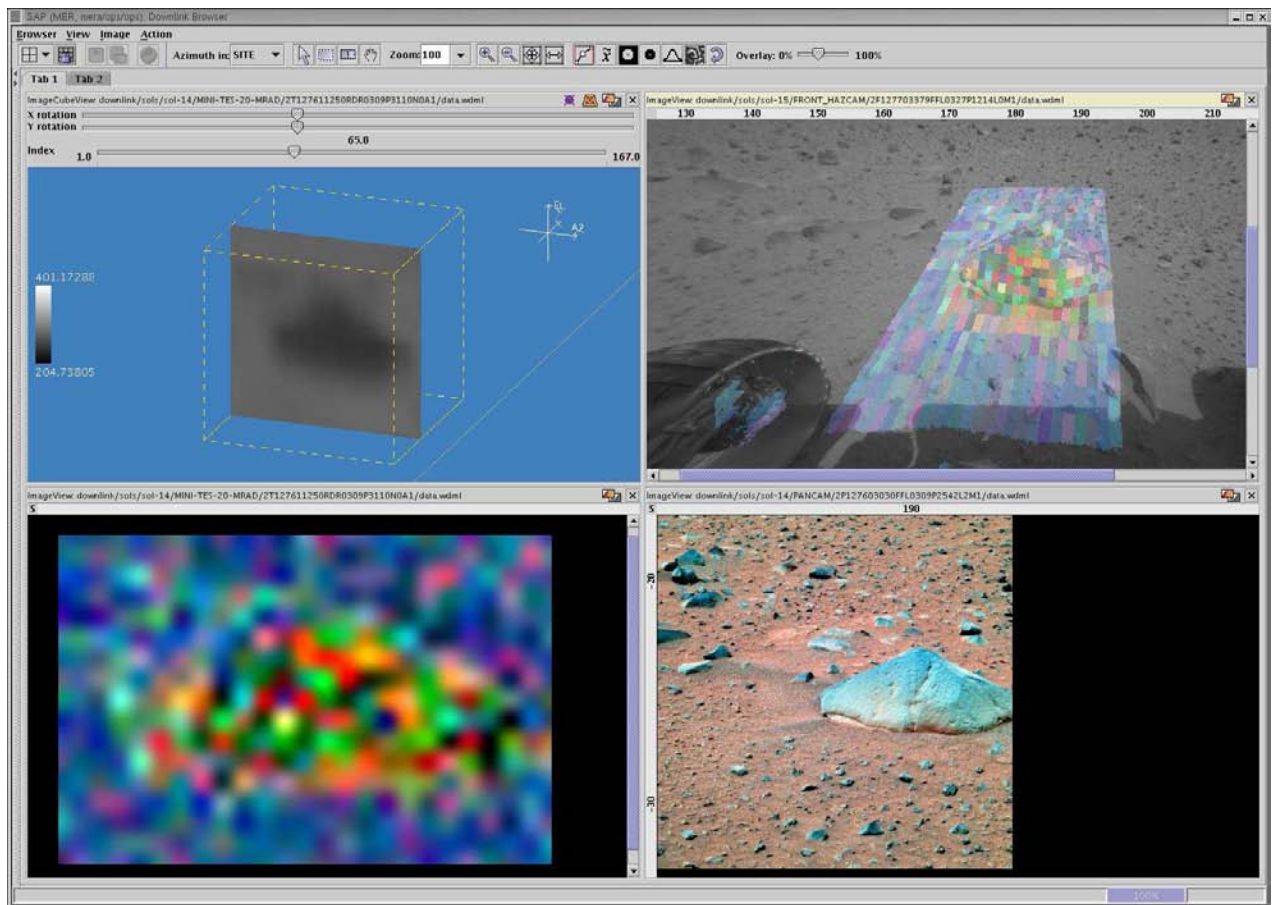


Fig. 5. Four views of the rock known as Adirondack at Gusev Crater. Upper left, one infrared spectral band observed by Mini-TES in an image cube view. Lower left, three infrared spectral bands combined into a synthetic RGB color image. Upper right, a data fusion of the synthetic color image from Mini-TES overlaid onto a front Hazcam image of the rock and surrounding soil. Lower right, a color Pancam image of this region.

of single images, cylindrical mosaics, and overhead maps were unanimously embraced, which speaks to the great versatility and accessibility of this style of visualization and high applicability in the presentation of imagery. For spectroscopy, image cube visualization was very useful for initial analysis of multispectral data. Data fusion was useful both for later stages of analysis (placing data of significant interest accurately in the context of stereo camera imagery) and for presentation of results to the science team as a whole, not all of whom are spectroscopy specialists. The 3D data visualization used most often by the scientists by far was anaglyph stereo visualization. Immersive 3D visualization was only used by a few scientists, although it was used extensively by those engineers who planned detailed traverse and arm placement activities. In summary, we feel that providing the extensive variety of scientific visualization capabilities reported herein was justified in that it served its expected role in serving the needs of the tactical science planning process for the Athena Science Team.

ACKNOWLEDGMENT

This work was performed for the Jet Propulsion Laboratory, California Institute of Technology, sponsored by the

National Aeronautics and Space Administration. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

REFERENCES

- [1] S. W. Squyres et al., "The spirit rover's athena science investigation at gusev crater, mars."
- [2] D. Gennery, "Camera calibration including lens distortion," Tech. Rep. JPL D-8580, Jet Propulsion Laboratory, California Institute of Technology, 1991.
- [3] D. B. Gennery, "Generalized camera calibration including fish-eye lenses," *unpublished*, 2002.
- [4] M. W. Powell, M. A. Vona, P. G. Backes, and J. S. Norris, "Visualization of coregistered imagery for remote surface operations," in *IEEE Aerospace Conference*, (Big Sky, MT), March 2003.
- [5] M. A. Vona and M. W. Powell, "Visible scalable terrain (vista) format."
- [6] M. A. Vona, P. G. Backes, J. S. Norris, and M. W. Powell, "Challenges in 3d visualization for mars exploration rover mission science planning," in *IEEE Aerospace Conference*, (Big Sky, MT), March 2003.
- [7] M. W. Powell, J. S. Norris, and P. G. Backes, "Visualization of spectroscopy for remote surface operations," in *IEEE Aerospace Conference*, (Big Sky, MT), March 2002.