**AUTONOMOUS SCIENCE ANALYSES OF DIGITAL IMAGES FOR MARS SAMPLE RETURN AND BEYOND.** V.C. Gulick[1], R.L. Morris[1], M. Ruzon[2] and T.L Roush[1], [1]NASA-Ames, MS 245-3, Moffett Field, CA 94035; [2]Stanford University, Dept. Computer Sciences, Stanford, CA; *email:* vgulick@mail.arc.nasa.gov.

**Introduction:** To adequately explore high priority landing sites, scientists require rovers with greater mobility [1]. Therefore, future Mars missions will involve rovers capable of traversing tens of kilometers (vs. tens of meters traversed by Mars Pathfinder's Sojourner). However, the current process by which scientists interact with a rover does not scale to such distances. A single science objective is achieved through many iterations of a basic command cycle [2]: (1) all data must be transmitted to Earth and analyzed; (2) from this data, new targets are selected and the necessary information from the appropriate instruments are requested; (3) new commands are then uplinked and executed by the spacecraft and (4) the resulting data are returned to Earth, starting the process again. Experience with rover tests on Earth [3] shows that this time intensive process cannot be substantially shortened given the limited data downlink bandwidth and command cycle opportunities of real missions. Sending complete multicolor panoramas at several waypoints, for example, is out of the question for a single downlink opportunity. As a result, long traverses requiring many science command cycles would likely require many weeks, months or even years, perhaps exceeding rover design life or other constraints.

Autonomous onboard science analyses can address these problems in two ways. First, it will allow the rover to transmit only "interesting" images, defined as those likely to have higher science content. Second, the rover will be able to anticipate future commands, for example acquiring and returning spectra of "interesting" rocks along with the images in which they were detected. Such approaches, coupled with appropriate navigational software, address both the data volume and command cycle bottlenecks that limit both rover mobility and science yield.

We are developing algorithms to enable such intelligent decision making by autonomous spacecraft. Reflecting the ultimate level of ability we aim for, this program has been dubbed the "Grad Student on Mars Project". We envision, for example, an appropriately intelligent Athena-like rover at the Pathfinder landing site might be able to traverse over the ridge towards "Twin Peaks" to obtain better information on the stratigraphy of these "streamlined islands" or of the size, composition and morphology of boulders located on them. Along the traverse, the intelligent rover would collect and analyze images and obtain spectra of geologically interesting features or regions. The intelligent rover might also traverse further up Ares Vallis, and find additional paleoflood stage indicators such as slackwater deposits. Recognizing additional regions where boulders are imbricated, noting changes in their size, distribution, morphology, composition and the associated changes in channel geometry would yield important information on the outflow channel's paleoflood history. Representative images and associated supporting data from these locations could be downlinked to Earth along with the data requested by scientists from the previous uplink opportunity.

Our initial work has focused on recognizing geologically interesting portions of images. Here we summarize some of the algorithms to date.

**Autonomous Image Analyses algorithms:**

*Rock Locator.* The ability to autonomously locate rocks and other geologic features of interest within an image would allow a rover to automatically downlink spectra of all rocks or only the *n* largest rocks within a given field of view along with the context image. Such an algorithm would obviate the need to wait several command cycles to obtain similar data. Our algorithm calculates the position of the sun at the time an image is taken, detects the shadow boundary formed by rocks, and estimates the image position as well as the size of the rocks causing the shadows. It also provides quick estimates of whether rocks are likely large enough to contain the spectrometer's field of view (FOV).

The ability to automatically select suitable targets in conjunction with automated spectral analysis could prove an invaluable asset for rover missions of increasing mobility. For example, a moving rover could periodically stop, acquire an image, select targets for spectral analysis, and analyze them for the presence or absence of carbonates or other interesting mineralogical features. As a second example, the rover could stop, acquire an image, take spectra of autonomously selected targets, and return the image along with the spectra, or a few bits or words representing the results of an on-board analysis, to Earth. If scientists are satisfied with the autonomously selected targets, then a command cycle is saved. If not, then spectra can still be acquired from the desired targets by manually designating the targets on the next uplink opportunity. The results of this algorithm can be used in conjunction with other algorithms described below, for example, to autonomously obtain textural information, spectral data, and high-resolution images of the segmented rocks or objects.

*Rock Segmentation.* We are developing an algorithm to find rocks in an image based on our rock locator. We use the boundary between a rock and its shadow as a "seed" to locate the entire rock boundary. Knowledge of rock size and shape, in addition to its location in the image, facilitates collection of other

relevant science data. The rover can determine the direction in which to point its spectrometer and decide whether the spectrometer's FOV is contained completely within the rock. Additional applications include taking high resolution images of rocks, as well as analyzing a rock's visual texture to determine which rocks in a scene differ from previous observations and might warrant follow up spectroscopy or close-up imaging. Although both the rock segmentation and rock locator algorithms potentially have several planetary applications, they may also be useful in terrestrial field studies. In collaboration with CMU, we plan to apply a version of these algorithms to locate and segment rocks in the identification of meteorites by a rover in Antarctica [4].

*Stratigraphic Layer detector.* We have been exploring ways of detecting stratigraphic layers within an image. Layers are interesting because they give information on geologic processes acting through time. Layers are produced at all scales by a variety of geologic processes, including fluvial, aeolian, volcanic, metamorphic, glacial and lacustrine processes. The ability to recognize them, particularly *in situ*, would potentially provide a powerful way in which to analyze the local geologic history. Our algorithm detects intensity edges and searches the image for regions in which the edges have a preferred orientation. The location of layers and their relative amount (as a percentage) are recorded. If layers can be identified within a scene, then further analyses, such as obtaining high-resolution images or spectral data of the layers, can be initiated. The amount of layering in an image can also be used as an "interest measure" for determining, for example, which 10 images from previous command cycle are the most relevant.

*Texture classifiers and image segmentation.* We have been exploring image segmentation through visual texture as a method of classifying different portions of an image. If an image contains a few boulders, many small rocks, and some hills in the distance, we would like to partition the image so that the boundaries of our partition match the boundaries between the different image parts. By creating local representations of texture (using e.g. Gabor filters, Markov Random Fields, or Fourier spectra) and grouping pixels with similar representations, we hope to achieve a "meaningful" segmentation. Afterwards, assuming that we are interested in, for example, the boulders, our result would allow the rover to compress regions not identified as boulders to a high degree, while regions containing boulders could be returned at lower compression rates. The ultimate, long term application of these algorithms, however, would be as an aid in mapping of geologic surface units. The algorithms would recognize new, interesting, or novel surfaces and insure that those images and associated spectra are returned along with other requested data at the next downlink opportunity.

*Sky Detection and Removal.* Removing the portion of an image containing the sky can reduce the size of images while preserving all essential geological information. We have developed fast algorithms that autonomously detect and remove the sky in images. These algorithms save bits previously devoted to returning sky and allow surface bits to be returned in either full resolution or at lower compression ratios than otherwise possible. Alternatively, in atmospheric studies where the sky is the region of scientific interest, just the sky pixels can be returned together with a buffer zone of surface pixels at the horizon for context. Any significant "bit" savings, allows additional images and other science data to be returned.

A simple algorithm removes a rectangular region of the sky a few rows of pixels above the highest land point in the image, assuring that no topography was inadvertently removed. In another algorithm, the image is cropped further, resulting in a non-rectangular image whose top border is always a few pixels above the horizon. Zeros are added, both to keep the image rectangular and to facilitate lossless compression algorithms such as run-length or LZW encoding. The algorithm removes more sky, resulting in greater savings.

In addition to reducing image size, the sky detector has several other applications. For example, knowledge about where the sky is in an image can be used to improve the efficiency of other autonomous science analyses algorithms or to refine camera pointing. If the image contains more sky than is desired, the sky detector can invoke a command to either not send back that image or to reposition the camera so that the image contains more of the surface.

**Conclusion.** Although these algorithms should be particularly useful to missions with a highly mobile rover, they should be useful for other missions (e.g., orbiters, airplanes, balloons, less capable rovers) that can acquire and process many images but can only return a small fraction due to bandwidth limitations. These algorithms are prototypes and will be improved based on feedback from scientists and performance during rover field tests [3]. We are planning to integrate these and other yet to be developed tools into a single system that more efficiently analyzes images in a single pass.

While autonomous image analyses will probably never replace the geologist, such innovations should enable planetary scientists in future Mars missions to obtain more science from a single command cycle and therefore explore the landing site and its surrounding environs more fully.

**References:** [1] Gulick, V.C., editor, (1998) Mars Surveyor Landing Site wrkshp abs. vol., Jan. 26-27, 1998, NASA-Ames, Moffett Field, CA or see web page at (http://cmex.arc.nasa.gov/Mars_2001/2001_abs_vol.html). [2] Roush, et al. (1999) Auton. Sci. ..., LPSC. [3] Stoker et al., (1999), LPSC. [4] Pedersen L. et al., (1998), LPSC.