AUTONOMOUS SCIENCE DECISION MAKING FOR MARS SAMPLE RETURN. T.L. Roush[1], V. Gulick[1], R. Morris[1], P. Gazis[1], G. Benedix[2], C. Glymour[3], J. Ramsey[3], L. Pedersen[4], M. Ruzon[5], W. Buntine[6], and J. Oliver[6]. [1]NASA Ames, MS 245-3, Moffett Field, CA 94035-1000 (troush@mail.arc.nasa.gov), [2]Dept. Geology, Arizona State Univ., Tempe, AZ 85287, [3]Carnegie Mellon Univ., Dept. of Philosophy, Pittsburgh, PA 15213, [4]Carnegie Mellon Univ., Robotics Inst., Pittsburgh, PA 15213, [5]Stanford Univ., Dept. of Computer Science, Stanford, CA 94305, [6]Ultimode Systems, Berkeley, CA 94704

In the near future NASA intends to explore Mars in preparation for a sample return mission using robotic devices such as landers, rovers, orbiters, airplanes, and/or balloons. Such platforms will likely carry imaging devices to characterize the surface morphology, and a variety of analytical instruments intended to evaluate the chemical and mineralogical nature of the environment(s) that they encounter. Historically, mission operations have involved the following sequence of activities: (1) return of scientific data from the vehicle; (2) evaluation of the data by space scientists; (3) recommendations of the scientists regarding future mission activity; (4) transmission of commands to the vehicle to achieve this activity; and (5) new activity by the vehicle in response to those commands. This is repeated for the duration of the mission, with command opportunities once or perhaps twice per day. In a rapidly changing environment, such as might be encountered by a rover traversing hundreds of meters a day or an airplane soaring over several hundred of kilometers, this traditional cycle of data evaluation and commands is not amenable to rapid long range traverses, discovery of novelty, or rapid response to any unanticipated situations. In addition, to issues of response time, the nature of imaging and/or spectroscopic devices are such that tremendous data volumes can be acquired, for example during a traverse. These data volumes can rapidly exceed on-board memory capabilities prior to an opprotunity to transmit it to Earth. Also communication band-widths are sufficiently restrictive that only a small portion of these data can actually be returned to Earth. These considerations clearly suggest that crucial decisions regarding data analysis will have to be made on-board by these robotic explorers. These decisions are distinct from, but are influenced by, the more conventional issues of electromechanical control, health, and navigation associated with robotic operations. Instead these decisions require automating scientific discovery by the robot craft based upon data returned by its sensors. Such an approach would eventually enable the robotic explorer to understand what is *interesting* because the data deviates from expectations generated by current theories/models of planetary processes that could be associated with the observed data. Such interesting data and/or conclusions can then be selectively transmitted to Earth thus reducing memory and communications demands.

At NASA Ames, planetary and computer scientists have joined forces in evaluating the framework and tools necessary to enable robotic craft to undertake such scientific exploration. Tools evaluated to date include: (1) Bayesian statistical classification of near-infrared spectroscopic and imaging data; (2) visual image segmentation and layer detection; (3) statistical clustering of near-infrared spectroscopic data; (4) neural network assays of mineral identities and relative abundances based upon near-infrared spectroscopic data; (5) automated Bayes net construction from statistical constraints, and (6) fusion of simultaneous image and spectral data guided by expert system knowledge obtained from geologists. A sub-set of these tools have recently been tested during rover operations in Antarctica and others will be tested in rover operations scheduled to begin in early 1999.

One autonomous scientific analysis to be evaluated is a test to determine the presence or absence of carbonate minerals based upon visual and near-infrared spectral reflectance data to be obtained during the rover field operations. This analysis relies upon two independent evaluations of the spectral data. The first analysis evaluates the characteristics of spectral features observed in the data, e.g. wavelength positions, widths, and relative strengths of absorption bands, and compares these values to those determined from spectral libraries. The second analysis relies upon the results of an unsupervised Bayesian classification scheme that used a spectral library to define a number of classes of minerals sharing similar spectral properties. A spectrum obtained in the field is evaluated to determine which of these predefined classes it belongs to. Each analysis alone may give rise to substantial uncertainty in a positive identification. For the first analysis method, this might occur because other minerals have spectral features that could be confused with the carbonates. For the second analysis method, this might occur because a variety of minerals may be combined in any one class determined by the Bayesian classification scheme. However, this uncertainty is reduced by combining the results of these analyses, e.g. a spectrum is required to have characteristics consistent with carbonates and the determined class must contain carbonates as a member. Such information could trigger other activity by the robotic explorer, e.g. obtain a sample of the carbonate, or to reduce data volume requirements. For example, the spectrometer intended for use in the robotic field operations requires 9000 bits to return an entire spectrum in a proprietary binary format whereas the results of the analysis described above could be contained in a single bit.

Other analysis techniques related to images involve both data volume reduction and scientific analysis/evaluation. One tool autonomously recognizes sky/horizon boundary and can reduce image data volumes by rejecting the geologically uninteresting sky regions from the images. Another technique involves autonomously recognizing contiguous locations within images that share similar properties. Such regions could contain layers, represent dust deposits between rocky outcrops, or be associated with individual rocks or boulders within an image. This information could be used to alter rover activity, e.g. obtain a spectrum of a layered deposit, image a specific region at higher spatial resolution, or move towards a specific region. Alternatively the information could be used to prioritize the order in which images are returned to Earth, e.g. images containing layers have the highest priority for transmission. Finally, such information could be used to reduce data volume requirements, e.g. potentially determine and return the histogram of a rock size distribution in an area instead of the image.

All of these tools require input from domain experts, testing under realistic conditions, and evaluation of their results. The field operations scheduled for early 1999 will provide an initial test for some of these systems under realistic conditions. The results of these tests will be evaluated in terms not only of data reduction, but more importantly in their effectiveness of conveying the scientific information contained in the original data.