Automatic Rock Detection and Mapping from HiRISE Imagery

NASA's Jet Propulsion Laboratory, Pasadena, California

This system includes a C-code software program and a set of MATLAB software tools for statistical analysis and rock distribution mapping. The major functions include rock detection and rock detection validation. The rock detection code has been evolved into a production tool that can be used by engineers and geologists with minor training.

The software takes as an input an image of a scene containing rocks and produces as output a description of the rock population and associated statistics. Each rock is described in terms of location, dimensions, and confidence of detection. The input parameters are the image resolution (ground sampling distance, or the size of a pixel in centimeters), the Sun incidence and azimuth angles for analysis of the shadows cast by rocks to derive individual rock models, and a parameter than can be adjusted to accommodate variations in image contrast.

The software is able to process very large reconnaissance imagery using a standard desktop computer by automatically processing image blocks and collecting all output in a single rock population description file (RPDF). Processing time is in the order of minutes for nominal HiRISE images covering 6×12 km areas at 30 cm/pixel.

The test option allows small portions of the large images to be selected and processed. Alternatively, a specific image window can be processed by indicating its coordinates and size. In this mode, visual results (detections overlaid on the images) are provided in addition to the rock population file (RPDF). This option is useful to quickly allow verification of parameter settings, and the quality of the detection results.

This work was done by Andres Huertas, Douglas S. Adams, and Yang Cheng of Caltech for NASA's Jet Propulsion Laboratory.

This software is available for commercial licensing. Please contact Karina Edmonds of the California Institute of Technology at (626) 395-2322. Refer to NPO-45752.

Parallel Computing for the Computed-Tomography Imaging Spectrometer

NASA's Jet Propulsion Laboratory, Pasadena, California

This software computes the tomographic reconstruction of spatial-spectral data from raw detector images of the Computed-Tomography Imaging Spectrometer (CTIS), which enables transient-level, multi-spectral imaging by capturing spatial and spectral information in a single snapshot. The CTIS can be used for surveying planetary landscapes through spectral imaging. It can also be used for battlefield surveillance and the spectral imaging of live tissues for disease detection.

A Message Passing Interface Library (MPI) is used to parallelize the original serial version of the code without modifying its initial structure. By parallelizing the code, a speedup of up to 20 is reached by using 32 processors. The software does not use any third-party libraries that require licenses. It is written in Fortran and MPI, and the storage of matrix elements is efficient, thus reducing memory requirements.

This work was done by Seungwon Lee of Caltech for NASA's Jet Propulsion Laboratory. Further information is contained in a TSP (see page 1).

This software is available for commercial licensing. Please contact Karina Edmonds of the California Institute of Technology at (626) 395-2322. Refer to NPO-45831.

Rock Segmentation Through Edge Regrouping

NASA's Jet Propulsion Laboratory, Pasadena, California

Rockster is an algorithm that automatically identifies the locations and boundaries of rocks imaged by the rover hazard cameras (hazcams), navigation cameras (navcams), or panoramic cameras (pancams). The software uses edge detection and edge regrouping to identify closed contours that separate the rocks from the background (see figure). The algorithm has applications both in ground-based data analysis, for example, to examine large quantities of images returned by the Mars Exploration Rovers, and in onboard (on-rover) opportunistic science applications such as construction of rock maps during traverse, identification of unusual or otherwise high-value science targets that warrant additional investigation, and detection of certain types of geologic contact zones.

The software uses gray-level intensity gradients to identify raw contours; these raw contours are then split into shorter, low-curvature fragments. New fragments are created where necessary to bridge areas of poor gradient information or poor image quality. The algorithm uses a flooding step to regroup the various fragments into closed contours. The algorithm is very fast with the C implementation able to process (768×1024) images containing hundreds to thousands of rocks in approximately one second on a desktop workstation.

The algorithm is particularly efficient at quickly detecting small- to medium-sized rocks with sufficient contrast (positive or negative) relative to the background. Full quantitative performance comparisons are not yet available; however, preliminary tests show that Rockster appears to detect a significantly larger fraction of rocks

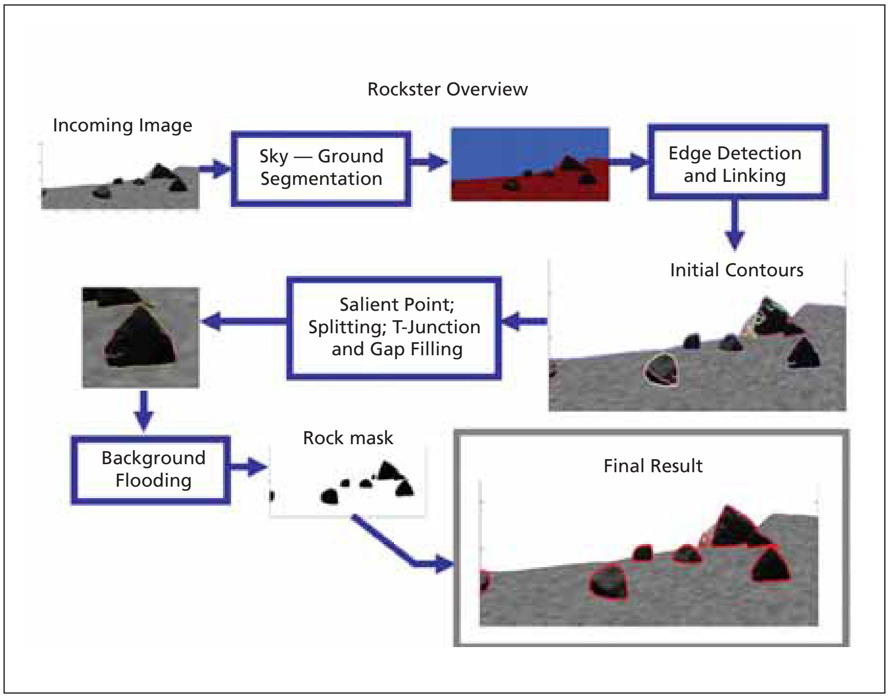

An **Overview of the Rockster** rock segmentation algorithm.

present in a scene (higher recall) than previous rock detection schemes, while maintaining a high precision rate (objects identified as rocks, truly are rocks).

Rockster has been integrated successfully into a number of recent, high-level demonstrations, including the SOOPS (Science Operations on Planetary Surfaces) demo, which used a rock exploration scenario to let scientists gain hands-on experience with an autonomous science capability in a simulated environment, and live exercises of the OASIS (Onboard Autonomous Science Investigation System)/CLARAty (Coupled Layer Architecture for Robotic Autonomy) software which were carried out in real-time in the JPL Mars Yard onboard the FIDO Rover (a close relative of the twin Mars Exploration Rovers).

This program was written by Michael Burl of Caltech for NASA's Jet Propulsion Laboratory.

This software is available for commercial licensing. Please contact Karina Edmonds of the California Institute of Technology at (626) 395-2322. Refer to NPO-44417.

System for Continuous Delivery of MODIS Imagery to Internet Mapping Applications

NASA's Jet Propulsion Laboratory, Pasadena, California

This software represents a complete, unsupervised processing chain that generates a continuously updating global image of the Earth from the most recent available MODIS Level 1B scenes.

The software constantly updates a global image of the Earth at 250 m per pixel. It uses an event-driven scheduler to manage asynchronous image generation tasks on a cluster of computers. The output composite image is tightly integrated in the JPL OnEarth WMS server, which offers direct access to the global image to any Web Mapping Service (WMS) compatible client, and also supports KML generation for Google Earth. The resulting Earth image composite is permanently available as an Internet service to WMS compatible mapping applications. This application can handle the throughput of the MODIS satellite, processing more than 80GB of Level 1B input data each day.

There are two main components to this software package: DailyHarvest and DailyPlanet. The first component is a scene harvester that manages a local copy of available MODIS scenes for the past few days. When active, the DailyHarvest module checks the current state of the MODIS source repository (LAADS) against the local state. Any new scene is downloaded, and the local copy state is altered to reflect the availability of the new scene. The second component, DailyPlanet, is then made aware of the change in the archive state. Once the remote and local scenes are synchronized, the DailyHarvest module reschedules itself, effectively running every few minutes.

The DailyPlanet module functions as an event-driven scheduler that manages the scene transitions from raw scenes to the global composite. There are three separate scene queues: a raw scene queue, a confirmed scene queue, and an image fragment queue. Each scene makes the transition to the next queue based on the result of an external process. The transition from raw-to-confirmed is done by extracting the scene metadata and applying a suitability test, confirming that the scene is not a night scene, and that the latitude range is within the S72-N72. The transition from the confirmed scene queue to the image fragment queue is the result of the successful completion of an external task that produces a visual image in geographical coordinates from a MODIS Level1B scene. This processing is done using the HDFLook MODIS software package, running on a remote computer.

The final task is the integration of the image fragments into the global composite, handled by a custom written external task. If an error is detected during any of the transitions, the error is logged and the scene is dropped from the processing queues. Most of the remote processing resources can be configured and are used in parallel for greater efficiency.

This work was done by Lucian Plesea of Caltech for NASA's Jet Propulsion Laboratory.

This software is available for commercial licensing. Please contact Karina Edmonds of the California Institute of Technology at (626) 395-2322. Refer to NPO-45778.