

# AUTONOMOUS ROVER REFLECTANCE SPECTROSCOPY WITH DOZENS OF TARGETS

Francisco Calderón<sup>1</sup>, David R. Thompson<sup>2</sup> and David Wettergreen<sup>3</sup>

The Robotics Institute, Carnegie Mellon University  
5000 Forbes Ave, Pittsburgh PA 15213 USA

<sup>1</sup>fc@alumni.cmu.edu, <sup>2</sup>drt@ri.cmu.edu, <sup>3</sup>dsw@ri.cmu.edu

## ABSTRACT

Future planetary rovers will travel autonomously over the visible horizon where data collection sequences cannot be scripted in advance. These rovers would benefit from an ability to autonomously identify science targets and deploy sensors. We present an automatic method for measuring rocks using a VIS/NIR reflectance spectrometer and solar illumination. Our approach leverages an automatic image-based rock detection algorithm and SIFT keypoints [7] that establish feature correspondence across images. This permits a rover to collect multiple spectral target measurements in a single command cycle with no direct human intervention. We present field tests of the system at Amboy Crater in the Mojave Desert. Here rock detection with visual servoing significantly improves the diversity and purity of collected spectra.

## 1. INTRODUCTION

Next-generation planetary exploration rovers will soon be able to travel autonomously over the visible horizon and visit multiple geologic units in a single command cycle. These surveys promise significant benefits for planetary geology but also raise an important challenge: at uplink scientists cannot know the specific features of interest that the robot will encounter. Data collection sequences cannot be scripted in advance, and the robot itself must autonomously identify science targets and deploy appropriate sensors.

We present an automated method to obtain spectroscopic measurements of rocks using a VIS/NIR reflectance spectrometer and solar illumination. Here a template-based detection algorithm finds rocks in each new image. Scale-Invariant Feature Transform (SIFT) keypoints [7] characterize previously-

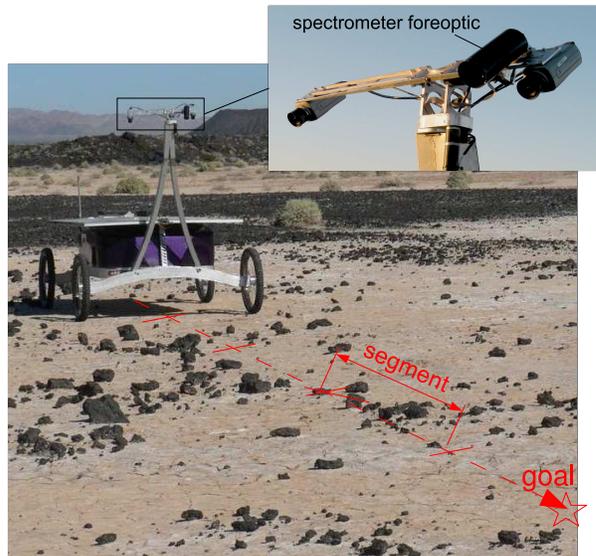


Figure 1. Zoë at the Amboy crater site. The rover travels between waypoints toward the goal while collecting spectra of rocks. Actual field tests use a traverse length of 50m. An image insert shows the stereo rig and spectrometer foreoptic.

appearing rocks. This correspondence matching yields a database with a unique position estimate for each science target. The robot periodically pauses its forward motion, performs visual servoing to aim the spectrometer, and collects measurements from the most promising targets. Our approach allows a robot to collect dozens or potentially hundreds of spectral measurements in a single command cycle with no direct human intervention.

Previous research has examined automatic spectroscopy in the context of Single-Cycle Instrument Placement (SCIP). This approach generally begins with a human specifying a target in a panoramic image. The rover advances toward the feature with

visual tracking, and then performs a “handoff” to instrument cameras for fine tracking and arm placement [1; 2; 8]. Pedersen et al. extend this model to multiple targets [9]; their system demonstrated high-accuracy instrument placement on 3 targets in under 3 hours. Previous studies have also considered the problem of autonomous rock detection. Castaño et al. show a system that incorporates rock detection with a resource-aware planner. Here the rover schedules opportunistic followup images of detected rocks that match a desired target profile. In other work, Thompson and Castaño’s survey of rock detection algorithms concludes that several are sufficiently accurate to facilitate autonomous spectroscopy [11].

Our work differs from previous efforts in several respects. First, we focus on reflectance spectroscopy under solar illumination which obviates the need for slow approach and arm placement maneuvers. Additionally, instead of reacting to a single prespecified target our algorithm autonomously detects and tracks multiple target rocks simultaneously. The autonomous spectrum collection has a relatively high failure rate compared to supervised SCIP methods. However, since there is no significant recovery penalty these errors are not especially dangerous. The rover can attempt many spectrum collection sequences and quickly acquire spectra from dozens of rocks. The end result of the procedure is a spectral map that provides a VIS/NIR profile of the transect.

This paper describes tests performed on a field rover at Amboy crater in the Mojave Desert, California (Figure 1). Here the rover traverses a field of rocks along the contact between an eroded basaltic lava flow and a sediment-covered plain. Accurate spectral pointing proves necessary to produce clean spectra that is uncontaminated by background sediment. Tests show that rock detection combined with visual servoing significantly improves the diversity and purity of collected spectra.

## 2. APPROACH AND ARCHITECTURE

We implement the spectral profiling algorithm on an exploration robot “Zoë” developed at Carnegie Mellon to test Mars-relevant autonomy technologies. Zoë’s science instruments include a 60cm-baseline stereo rig and an ASD Fieldspec Pro VIS/NIR 350 – 2500nm reflectance spectrometer. The spectrometer objective lens, or *foreoptic*, is mounted with the cameras on a pan-tilt base that provides full 360° coverage of the environment. The foreoptic’s field of view forms a 1° sensing cone in space. Placing a target into this cone allows the spectrometer to collect reflectance spectra under solar illumination.

In our operational procedure an operator specifies a transect using a start location, a goal waypoint, and a time budget. The rover drives from the start loca-

tion to the goal, producing a map of detected rocks and spectra — that is, a *spectral profile* of the rocks. The mapping procedure enforces a time constraint; one could integrate it as an atomic command into a larger mission plan with other drive or sampling actions.

Figure 2 shows a software architecture developed for the spectral profiling task. Our architecture assigns the tasks of data acquisition, image analysis, and navigation to three independently-operating software modules.

During operation the rover alternates between two states: a *travel* state for navigation and a *pause* state where it stops to collect spectra. During the travel state the rover follows a series of waypoints between the start location and the goal (Figure 1). The navigation module drives the robot from one waypoint to the next; it performs obstacle avoidance using stereo terrain analysis and a D\* path planning algorithm. Meanwhile the data collection module commands forward-looking images. The data analysis module examines these images and inserts newly detected rocks into a growing database of features.

A switch to the *pause* state requires two conditions: First, there must be at least one unmeasured rock within spectrometer range; and second, the current traverse segment must retain sufficient time for a spectrum collection sequence (to be described in Section 4). The navigation module stops the rover whenever these criteria are met. This signals the data collection module to begin spectrum collection. Every successful sequence results in a new spectral measurement which is then inserted into the database and associated with the correct rock.

The rock database is central to the system and acts as a medium by which the other components communicate. We implement this structure with a relational database. The following information is associated with each rock:

- An integer used as a globally unique identifier.
- A global 3D position estimate given in ECEF coordinates.
- Pixel bounding boxes for each image in which the rock appears.
- All SIFT features associated with the rock in previous images.
- Visual and spectral attributes of the rock.

Details of image processing (rock detection, visual features and rock correspondence) and spectrometer pointing appear in Chapters 3 and 4, respectively.

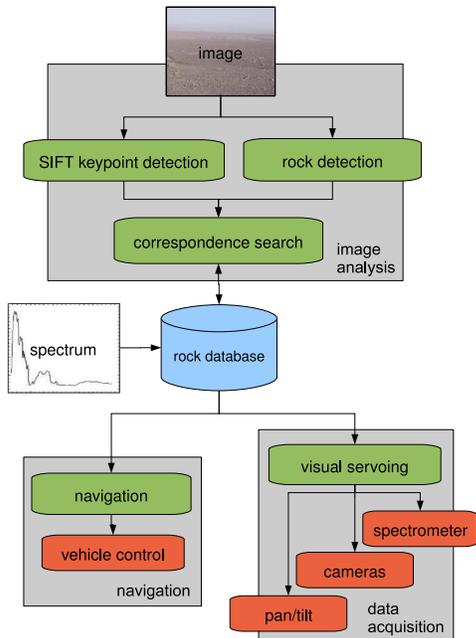


Figure 2. General architecture.

### 3. IMAGE ANALYSIS: ROCK DETECTION AND CORRESPONDENCE

The system detects rocks using a variant of the Viola/Jones approach [12]. This method identifies rock bounding boxes with a cascade of Haar-wavelet-like filters. An Adaboost supervised learning routine identifies a cascade of filters that discriminates between bounding boxes that contain rocks and those that do not. Applying this cascade to a novel candidate bounding box yields a real-time detection algorithm with reasonable rock detection performance.

In the rock detection domain, the lighting angle significantly affects rocks’ appearance. We compensate by training several parallel detectors on rocks lit from different angles as described in Thompson et al. [11]. For a novel image we apply all cascades and use only the results from the one returning the most detections. The rationale is that patches matching the learned appearance of rocks are unlikely to occur by chance. Most images contain several true rocks, so the cascade which detects the most rocks usually matches the illumination of the scene.

The image analysis stage filters detected rocks less than 10cm in size that are too small for accurate targeting. In addition we eliminate rocks that are within 2m of the rover center; this prevents spurious detections from the rover body or shadow from entering the database. The filtered result is a set of valid potential spectrometer targets; Figure 3 shows an example rock detection result from the Amboy Crater tests.

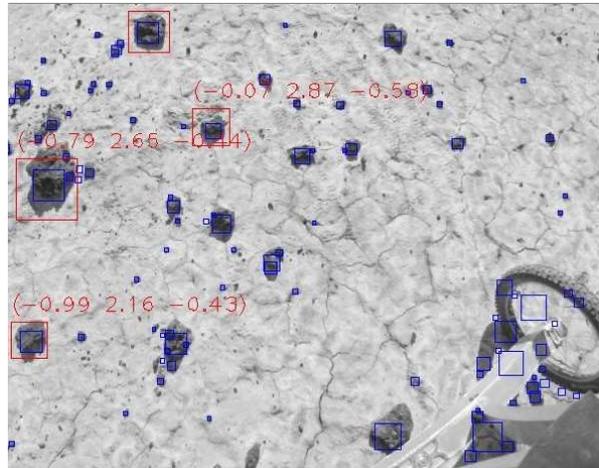


Figure 3. Detection result after filtering by rock size and distance. Large red bounding boxes are potential spectrometer targets, with numbers indicating rover-relative position estimates. Small blue squares indicate SIFT keypoints that are used for the correspondence search. The front axle is visible in the lower-right corner of the image.

#### 3.1. Rock Correspondence

SIFT keypoints are stable image locations that are invariant to scale and rotation and partially invariant to affine distortion, change in 3-D viewpoint, addition of noise and changes in illumination [7]. The image analysis software associates each rock to SIFT keypoints in its bounding box. We establish correspondence across images by matching new rocks’ SIFT keypoints against previous keypoints stored in the database.

A single traverse might view thousands of unique keypoints so an unconstrained  $O(n^2)$  correspondence search is unsuitable for real-time visual servoing. We constrain the search with geometric information. Given the known absolute pose of the rover and the relative orientation of the cameras the system determines an approximate region (a radius of several meters expressed in ECEF coordinates) expected to contain the new SIFT keypoints. Keypoints that lie outside this region are excluded from the correspondence search.

We compare SIFT keypoints using the standard 128-dimensional descriptor vectors describing the local scale-invariant appearance of the patch [7]. We evaluate correspondence between SIFT descriptors by calculating the Euclidean distance of the incoming descriptor against each of the remaining candidates in the database. We use the traditional SIFT match criterion that requires a nearest-neighbor distance at least  $n$  times smaller than the distance to the second closest neighbor for an empirically-defined threshold  $n$  [7].

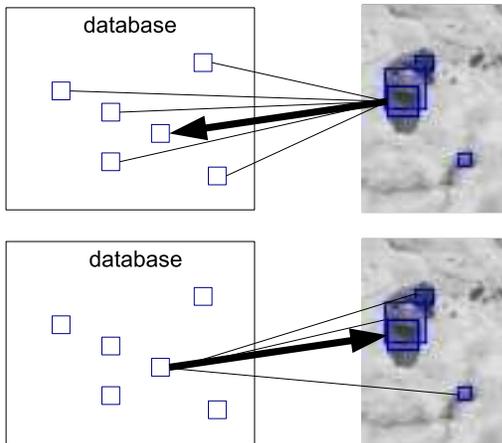


Figure 4. Correspondence with an existing SIFT descriptor in the database requires that the match be valid in both directions.

An additional criterion further reduces the number of false SIFT matches. For each of the matches from the Euclidean search we do an equivalent “backward” search over SIFT descriptors in the new image. In other words, we apply a similar Euclidean search for nearest and second-nearest neighbors using the old SIFT descriptor as the query and the new image’s SIFT descriptors as the search space (Figure 4). A match is only accepted when the thresholded Euclidean match is valid in both directions.

A successful match implies that the SIFT descriptor lies on a previously-detected feature. We update the rock’s visual and position information along with the SIFT descriptor’s record to reflect the fact that illumination and aspect may have changed. We also merge all instances of the rock in the database to eliminate duplicate entries. If none of a new rock’s SIFT descriptors match previous features it receives a new unique entry in the feature database. SIFT keypoints with no matches and that are not identified as part of any rock are assumed to belong to the background and discarded.

Note again that we perform the correspondence search on all SIFT keypoints in the incoming image, not just those belonging to a rock detection. In general rock detection performance is less robust than SIFT keypoint extraction; the rock detector often fails to re-acquire a rock after changes in the viewpoint. Therefore all SIFT features in the incoming image are potential matching candidates. In order to avoid flooding the database with irrelevant “background” keypoints we only record those keypoints that match existing descriptors or belong to a newly detected rock. In other words, all SIFT descriptors stored in the database must have been identified as part of a rock at least in one image.

### 3.2. Position Estimation

Two redundant methods estimate rock positions: stereopsis and the groundplane. Stereopsis relies on a 3D reconstruction based on SIFT keypoints. The robust matching algorithm (detailed above in Section 3.1) identifies any SIFT keypoints that appear in both left and right cameras. A least-squares geometric reconstruction using a pinhole camera model and radial distortion [4] yields a sparse reconstruction. We use the 3D centroid of matched SIFT keypoints as the rock’s position estimate. Alternatively, the *groundplane* position estimation method serves as a fallback whenever stereo matches fail. Groundplane estimation uses a single image and assumes the rock lies on a planar surface. This may be incorrect so we favor stereopsis whenever SIFT matches are available.

## 4. DATA ACQUISITION: SPECTROMETER POINTING AND ANALYSIS

The spectrometer pointing procedure attempts to aim the spectrometer foreoptic at a target given the target’s spatial position. This sequence has two stages. The first *kinematic* step leverages a kinematic model of the the pan/tilt mount and the most recent rock position estimates from the database. The second *feedback* step refines the initial pointing with an iterative visual servoing approach that places the target in a specific region in the image.

### 4.1. Calibration Procedure

Both kinematic and feedback steps rely on generic models representing the location of the foreoptic cone. We expect to place the center of the cone on the target to maximize the probability of covering the sample in full. We fit the parameters of the physical model using calibration data obtained in laboratory tests. The main challenge of this calibration is to accurately determine where the spectrometer is pointing. Since there is no visual trace to determine the field of view it is not possible to get 3-D measurements directly. Instead we infer the field of view using spectroscopic measurements of a known target. Our procedure uses a commercial 532nm green laser pointer as a reference. For practical purposes the narrow-beam laser projects to a single point in space. The laser spot is easy to detect in collected spectra due to a high peak in its center frequency.

### 4.2. Kinematic Model

Calibration points produce a kinematic model of the foreoptic on the robot. Generic forward kinematic

equations solve for the pan and tilt angles  $\theta$  and  $\alpha$ :

$$\begin{aligned}\theta &= f_{\theta}(x, y, z, h, d, \beta) \\ \alpha &= f_{\alpha}(x, y, z, h, d, \beta)\end{aligned}\quad (1)$$

Here  $(x, y, z)$  is the location of the rock in 3-D space,  $h$  is the height of the robot,  $d$  is the distance between the axis of rotation and the foreoptic, and  $\beta$  represents other less-relevant parameters such as rotations and translations of the foreoptic with respect to the base. We perform an unconstrained nonlinear optimization over the parameters trying to minimize the total square error of the predicted position. The result is an estimate for  $h$ ,  $d$  and  $\beta$ . We refer the reader to work by Calderón for additional details on the kinematic model and calibration procedure [3].

### 4.3. Visual Correction

It is impractical to maintain a sufficiently accurate model because many calibration parameters are unstable or vary with time. Examples are the changing tire pressure on the vehicle, inaccurate pan/tilt angle readings, and occasional replacement of parts. Therefore a spectrometer pointing system based solely on a feedforward kinematic model will inevitably exhibit some pointing error. We incorporate a second closed-loop visual correction step to correct these errors and improve pointing accuracy.

Our visual correction method exploits the fixed transform between the camera and the foreoptic. Both devices are mounted to a single platform so their relative position is constant regardless of the mobile base's pan and tilt (Figure 1 inset). Therefore the center line of the spectrometer sensing cone (*cc line*) is perceived as static by the camera and projects onto a line in the image. The intersection of the target with this line is fully determined by the distance from the foreoptic to the target.

The *cc line* is not visible in imagery but its position can be inferred using the laser reference. Since the laser point is centered in the foreoptic's field of view, the laser spot in each calibration image corresponds to a visible point on the *cc line*. We can reconstruct the entire *cc line* from multiple laser points at different ranges. Perspective projection suggests a model relating the depth of the target and associated row/column pixels (Figure 5).

Pointing the spectrometer at a target is then equivalent to varying the position of the foreoptic until the target appears in a specific pixel of the image. This defines an "error" in visual space based on the current and desired position of the target. The system applies a simple visual servo [5] to correct it. We determine the *Image Jacobian* that relates changes in the camera position to changes in the image. This

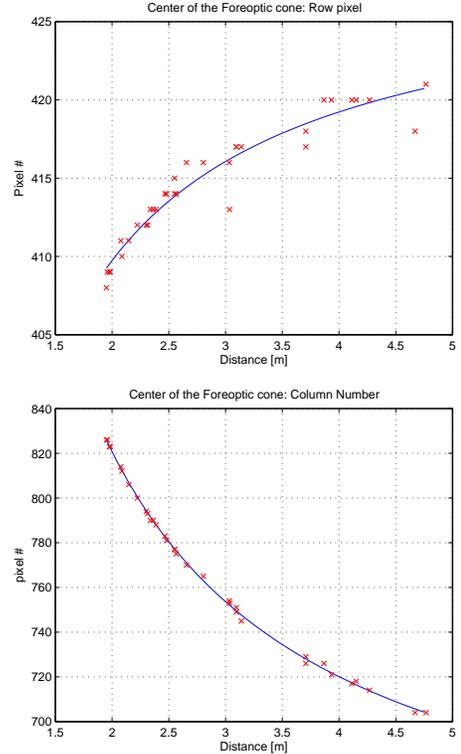


Figure 5. Row and column pixels of the foreoptic center in the image given different ranges. Charts show empirical calibration points and the best-fit model.

permits a proportional visual control law:

$$\begin{bmatrix} \omega_{tilt} \\ \omega_{pan} \end{bmatrix} = K J_{\nu}^{-1} \begin{bmatrix} \Delta u \\ \Delta v \end{bmatrix}\quad (2)$$

In this equation  $\omega_{pan}$  and  $\omega_{tilt}$  are the estimated corrections in pan and tilt, respectively.  $K$  is a constant gain matrix of the control law,  $J_{\nu}$  is the image Jacobian, and  $(\Delta u, \Delta v)$  is the visual error in both image axes. The procedure iterates the visual correction until the error disappears. The result is a system that can point accurately target features at ranges up to 5 meters.

### 4.4. Collection of White References

The rover periodically renormalizes the spectrometer against a white reference target. This compensates for changes in illumination during the 40 minute traverse due to sun angle and atmospheric phenomena. We employ a standard VIS/NIR white reference target exhibiting high diffuse, Lambertian reflectance that is mounted to the rover deck. The rover collects a spectrum from the reference target every 7 minutes using a scripted sequence of pan-tilt actions. All other software modules pause for this collection procedure. During the renormalization the rover straightens its axles and servos the pan-tilt unit to

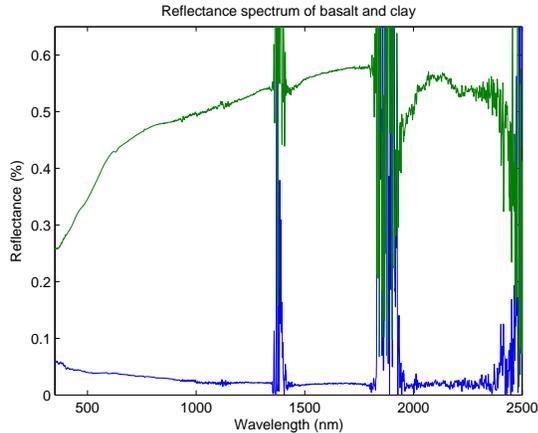


Figure 6. Spectra of basalt (blue - bottom line) and clay (green - top line) at Amboy Crater under solar illumination. High-noise ranges are caused by water absorption bands in the atmosphere.

target the white reference. Rough terrain sometimes results in poorly-oriented axles and — since the mast is fixed to the front axle — incorrect pointing. The rover senses this condition and aborts the attempt in favor of another 60 seconds later. Occasionally several attempts fail in succession; in these circumstances the rover continues trying every 60 seconds until a white reference spectrum is successfully captured. After a white reference the rover resets the standard 7-minute clock.

## 5. EXPERIMENTAL RESULTS

Experiments tested system performance in natural terrain. We transported Zoë to a site at Amboy Crater, California where two distinct materials can be distinguished both visually and spectroscopically. Basalt from lava flows appears as a dark material, while the underlying clay is lighter in color. This difference is also clear in VIS/NIR reflectance spectra (Figure 6).

We performed experiments over consecutive days in the fall of 2007. Clay predominates at the traverse site, but scattered basaltic rocks up to 30cm in length are also visible. Given that only these two materials are present, the collected spectra easily quantify the system’s rock targeting accuracy. We performed 4 experiments at different times of day under varying atmospheric conditions. The robot traveled approximately the same 50m transect with each trial.

Figure 7 provides the spectral profile resulting from a typical transect. The background shows a georegistered visible-band overflight image for context. Small black dots show rock detections from the database, while large colored dots illustrate collected spectra.

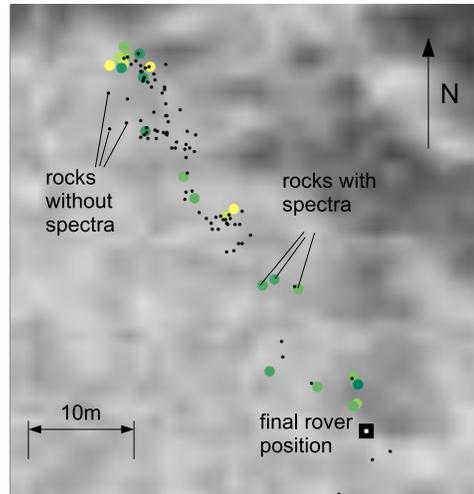


Figure 7. A spectrometer profile of a transect. A white square shows the final rover position. Small black and large colored dots respectively represent detected rocks with and without associated spectra. The map reflects the fact that density of rocks is high at start of the traverse and decreases near the end.

Table 1 gives performance results from each trial; the following sections explain each column in detail.

### 5.1. Rock Detection

We use the rock detection success criterion of Thompson et. al [10; 11] requiring that at least 50% of the detection contain a rock. The fourth column of Table 1 summarizes the resulting detection scores. The accuracy of our system is consistent with, or slightly better than, previously reported performance for this rock detection algorithm on benchmark datasets [11]. We attribute any performance advantage to particularly favorable field conditions: the images all show dark, regularly-shaped rocks against an even light-colored background. Note that detection performance decreases at later hours of the day as the zenith angle of the sun increases. Long shadows are generally absent from the training set and often cause false-positive detections.

### 5.2. Tracking and Correspondence

Here we consider system performance in completing the visual servoing procedure. The main failure mode occurs when the system is unable to match a target’s previous SIFT descriptors to those in a new visual feedback image. For a successful correction sequence the system must match some new SIFT descriptor to previous descriptors, finding the intended rock and reacquiring it for each of the 2 – 5 visual servoing iterations. Overall sequence success rates

Table 1. Details of individual trials at Amboy crater.

Trial	Time	Conditions	Rock Detect. Precision	Tracking Successes	Basalt Spectra
1	12:26 - 13:06	Moderate Clouds	332/361 (92.0%)	26/50 (52.0%)	22/26 (84.6%)
2	13:15 - 13:55	Sun	249/280 (88.9%)	32/52 (61.5%)	27/33 (81.8%)
3	11:38 - 12:18	Sun	280/296 (94.6%)	29/43 (67.4%)	18/29 (62.1%)
4	12:27 - 13:07	Sun	299/340 (87.9%)	26/46 (56.5%)	17/26 (65.4%)

appear in the fifth column of Table 1. This measure does not consider the feature itself (i.e. whether the target was a true rock or a false detection), but is based simply on the system’s ability to track arbitrary SIFT descriptors. Thus, it eliminates the influence of the rock detector.

Many descriptors stored in the database never reappear during the visual correction process so it is not always possible to reacquire a particular rock. Several causes may be at fault: changes of appearance of SIFT keypoints due to changes in illumination, rover pose, or shadows of the rock or the rover; a large number of similar SIFT descriptors in the database that prevents any confident match; or a SIFT keypoint that lies outside the camera’s field of view (unlikely in practice due to wide-angle cameras and robust rover position estimates).

Analyzing individual trials, we observe a lower performance in the first trial than the second (trials occurred consecutively on the same day). Many clouds were passing overhead during the execution of the first trial, causing significant changes in the light levels and complicating the SIFT descriptor matching. This is reflected in the captured data: failures cluster in adjacent tracking attempts whenever environmental conditions change abruptly. Trials 3 and 4 were also taken consecutively. The success tracking rate for trial 3 is slightly better; this may reflect a slower rate of illumination change when the zenith angle of the sun is higher.

In general we observe a tracking performance of over 55% when the light levels are favorable. An increased number of tracking failures causes the system to spend more time looking for known rocks present in the database, but does not affect the quality of collected spectra.

### 5.3. Spectrometer Pointing

Finally we quantify the accuracy of the spectrometer pointer to physically track rocks. We manually analyze each captured spectrum and classify it either as basalt or clay given that the spectra for both materials are known (Figure 6). The final column of Table 1 summarizes the result.

Classification of signals as soil or basalt is based on a principal component analysis (PCA) noise-reduction strategy, followed by a simple linear deconvolution. In principal component space only one variable exhibits a large variance; this corresponds to the basalt/clay distinction. We use PCA as a noise-reduction strategy, projecting collected spectra onto the first component vector and reconstructing the signal to produce a filtered spectrum. In order to interpret the PCA coefficient we apply a linear deconvolution with representative samples of pure basalt and clay. The result is the percentage of basalt that is present in the mix. Spectra showing at least a 50% basalt are classified as a success.

The system presents a relatively high variance when hitting targets but consistently exhibits accuracies higher than 62%. Many factors influence the number of missed targets; some are intrinsic to the method while others relate to environmental conditions. This causes changes between trials and a large between-trial variance. Factors affecting system accuracy include: stereo estimate errors; errors introduced by the ground plane assumption when stereo is not available; spurious rock detections; tracking visual features lying out of the rock (e.g. background or shadow); tracking visual features lying in the border of the rock (SIFT features are not guaranteed to be in the center of the rock); and SIFT mismatches during the tracking procedure.

### 5.4. Comparative Results

We compare system performance against *status quo* capabilities by performing three additional trials in the same testing area. In these trials the robot holds the pan/tilt unit at a fixed  $-30^\circ$  declination and blindly captures spectra at regular intervals during forward travel. Slight, inevitable perturbations of the start location (on the order of 10 – 20cm) ensure that the foreoptic’s field of view sweeps over different areas for each 50m traverse.

The results of the blind spectral mapping trials appear in Table 2; no spectral measurement of basalt appears in any of the spectra from any of the trials. This demonstrates the difficulty of the testing area; rocks were generally small and scattered. Therefore

Table 2. Tracking performance for periodic (blind) spectroscopy.

Trial	Basalt Spectra
1	0/68 (0.0%)
2	0/67 (0.0%)
3	0/64 (0.0%)

it is highly unlikely that a target would fall by chance into the spectrometer field of view during the periodic spectrum acquisition. Our pointing strategy outperforms the blind method; each trial measures at least at 26 individual rocks over a 50m traverse lasting no longer than 40 minutes.

## 6. CONCLUSIONS

This work describes the development, implementation and testing of new techniques for automatic reflectance spectroscopy for multiple rocks in natural terrain. The system can autonomously travel long distances, detect rocks and acquire spectral measurements on the fly. These results are relevant to science autonomy for next-generation planetary exploration systems that accumulate spectral databases and build geologic maps of the traversed regions. The main contributions of the work are the following:

- A system for rover-based VIS/NIR reflectance spectroscopy using solar illumination and an automatic white-reference calibration.
- The integration of autonomous rock detection with correspondence matching through SIFT features. Together these permit a database of unique rocks and their positions.
- Development of a procedure to calibrate a spectrometer to rover frame and camera views based on a kinematic targeting followed by a visual feedback correction. The procedure is capable of targeting individual features at ranges of up to 5 meters.
- A new data product — the *spectral profile* — that allows the piecewise integration of autonomous spectra collection into mixed mission plans containing both scripted and autonomous actions.

We have tested this system in a planetary analog field setting and measured its performance in multiple trials. In each trial the rover traverses 50 meters in 40 minutes while consistently collecting spectra of at least at 26 individual rocks.

## ACKNOWLEDGMENTS

The authors thank Dom Jonak and James Teza. We also thank Ron Greeley, Phil Christensen, and Shelby Cave for their help in choosing and understanding the Amboy field site. Finally we thank Analytical Spectral Devices, Inc. for their technical assistance. This work was supported by NASA ASTEP grant NNG0-4GB66G, “Science on the Fly.”

## REFERENCES

- [1] P. Backes et al., “Automated Rover Positioning and Instrument Placement,” *IEEE Aerospace* 2005.
- [2] M. Bajracharya, et al, “Target Tracking, Approach, and Camera Handoff for Automated Instrument Placement,” *IEEE Aerospace* 2005.
- [3] F. Calderón, ”Autonomous Reflectance Spectroscopy by a Mobile Robot for Mineralogical Characterization,” *Masters Thesis, The Robotics Institute, Carnegie Mellon University.* (publication pending) 2008.
- [4] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd Ed. Cambridge Univ. P., 2004.
- [5] S. Hutchinson, G.D.Hager and P.I.Corke, ”A tutorial on visual servo control”, *IEEE Transactions on Robotics and Automation*, 12(5):651-670, 1996.
- [6] T. Huntsburger, et al, “Closed Loop Control for Autonomous Approach and Placement of Science Instruments by Planetary Rovers,” *IROS* 2005.
- [7] D. Lowe, ”Distinctive image features from scale-invariant keypoints”, *International Journal of Computer Vision* 20:91-110, 2003.
- [8] R. Madison, “Improved target handoff for single cycle instrument placement.” *IEEE Aerospace* 2006.
- [9] L. Pedersen, M. Deans, D. Lees, S. Rajagoplan, D. E. Smith, ”Multiple Target Single Cycle Instrument Placement,” *ISAIRAS* 2005.
- [10] D. R. Thompson and S. Niekum and T. Smith and D. Wettergreen, R. Castaño, ”Automatic Detection and Classification of Geologic Features of Interest”, *IEEE Aerospace* 2005.
- [11] D. R. Thompson and R. Castaño, ”A performance comparison of rock detection algorithms for autonomous planetary geology”, *IEEE Aerospace* 2007.
- [12] P. Viola and M. Jones, “Rapid object detection using a boosted cascade of simple features.” *CVPR* 2001.