Science Autonomy for Rover Subsurface Exploration of the Atacama Desert

*David Wettergreen, Greydon Foil,
Michael Furlong, David R. Thompson*

■ *As planetary rovers expand their capabilities, traveling longer distances, deploying complex tools, and collecting voluminous scientific data, the requirements for intelligent guidance and control also grow. This, coupled with limited bandwidth and latencies, motivates on-board autonomy that ensures the quality of the science data return. Increasing quality of the data requires better sample selection, data validation, and data reduction. Robotic studies in Mars-like desert terrain have advanced autonomy for long-distance exploration and seeded technologies for planetary rover missions. In these field experiments the remote science team uses a novel control strategy that intersperses preplanned activities with autonomous decision making. The robot performs automatic data collection, interpretation, and response at multiple spatial scales. Specific capabilities include instrument calibration, visual targeting of selected features, an on-board database of collected data, and a long-range path planner that guides the robot using analysis of current surface and prior satellite data. Field experiments in the Atacama Desert of Chile over the past decade demonstrate these capabilities and illustrate current challenges and future directions.*

Robotic explorers communicate only intermittently with scientists because of limited opportunities for visibility by Earth-based antennas and the growing number of spacecraft needing attention. The data rate of deep space communication is also very limited. Autonomy can significantly improve science productivity in intervals between communication opportunities. In particular, science autonomy employs on-board analysis to make decisions affecting the scientific measurements that will be collected or transmitted.

We define *science autonomy* as using information about science objectives and interpretation of science instrument data to determine rover actions. Science autonomy encompasses detection and intelligent selection of measurements and samples, automatic acquisition of measurements. This includes automated approach and instrument/tool placement as well as calibration and verification, meaning collecting the intended measurement or sample. Intelligent collection of scientific measurements can increase both the quantity and quality of information gathered. Science autonomy describes utilizing scientific information to guide rover actions, for example, to execute an intelligent survey

Figure 1. Zoë in the Atacama Desert.

or mapping strategy that adapts as data is collected. Decisions about which samples to acquire and where and when to travel next can be based upon metrics of information gain. Similar metrics can also be used to prioritize science data for download. Intelligent compression strategies use knowledge or models of content to interpret and summarize in a compact form. The ultimate goal of science autonomy is to embody sufficient understanding, quantified by models and metrics, so that rovers can independently choose actions that best support the scientific investigation in which they are engaged. Rovers will take their goals and guidance from scientists, but when isolated they should make scientifically rational decisions and when in communication they should provide the most relevant information possible.

Science autonomy is especially valuable for surface rover operations because missions have finite lifetime and rarely revisit sites after the first encounter — the rover must make good decisions and get it right the first time. Recent demonstrations on spacecraft show increasingly sophisticated science autonomy capabilities. Milestones include target tracking during the

Deep Impact comet flyby (Mastrodemos, Kubitschek, and Synnott 2005); target detection and response by the Mars Exploration Rovers (Castaño et al. 2008; Estlin et al. 2012); and spectral detection, discovery, and mapping by the EO-1 spacecraft (Chien et al. 2005; Davies et al. 2006; Doggett et al. 2006; Ip et al. 2006; Thompson et al. 2013). At the same time, new smart instruments are beginning to incorporate autonomous science data analysis directly (Wagstaff et al. 2013) and provide information that can be used to guide the rovers' targeting and operation.

These techniques and others will enable surface rovers to achieve multiday autonomous operations. Currently multiday rover plans do not travel over the horizon of yesterday's imagery, which limits the daily science yield. However, rover navigation already permits safe over-the-horizon traverses, and in principle a rover could autonomously survey large areas of terrain with its full suite of instruments. In one natural arrangement, operators would direct the rover using waypoints determined from satellite images, relying on rover autonomy for low-level hazard avoidance and science target selection en route.

Figure 2. Locales Visited During the 2013 LITA Field Season.

A robot could even divert its path slightly to pursue science targets of opportunity (Woods et al. 2009). Multiday plans could therefore make very efficient use of communications and personnel resources, enhancing long-distance survey missions.

The Life in the Atacama project is a NASA-sponsored effort to evaluate these techniques in the context of desert subsurface biogeology (Cabrol et al. 2007). It uses Zoë (Wettergreen et al. 2008), a rover capable of traveling more than 10 kilometers per day and autonomously drilling up to 0.7 meter depth (figure 1). As a mobility platform it combines navigational autonomy with a changing payload of onboard science instruments. Previous investigations have used a fluorescence imager capable of detecting very specific organic compounds and neutron detectors to measure hydrogen abundance. The current configuration incorporates a Raman spectrometer, a visible near infrared point spectrometer, and navigation and science cameras. During a series of experiments in the summer of 2013, scientists guided Zoë

remotely through the desert while exploring its geology and biology.

This article describes the science autonomy system developed and tested with Zoë. It performed automatic acquisition of visible/near infrared (Vis-NIR) reflectance spectroscopy throughout the 2013 field season. This involved a range of different autonomous decisions exercised at various spatiotemporal scales. We begin by describing the rover platform and instrument payload. We then discuss instrument self-calibration, science feature detection, and targeting capabilities. We describe larger-scale path planning used to select informative paths between waypoints. We also detail the operational protocols used to command the rover and the results of its autonomous data collection. These experiments provide a case study of science autonomy deployed continuously over long distances. We report on system performance, lessons learned, and plans for future development.
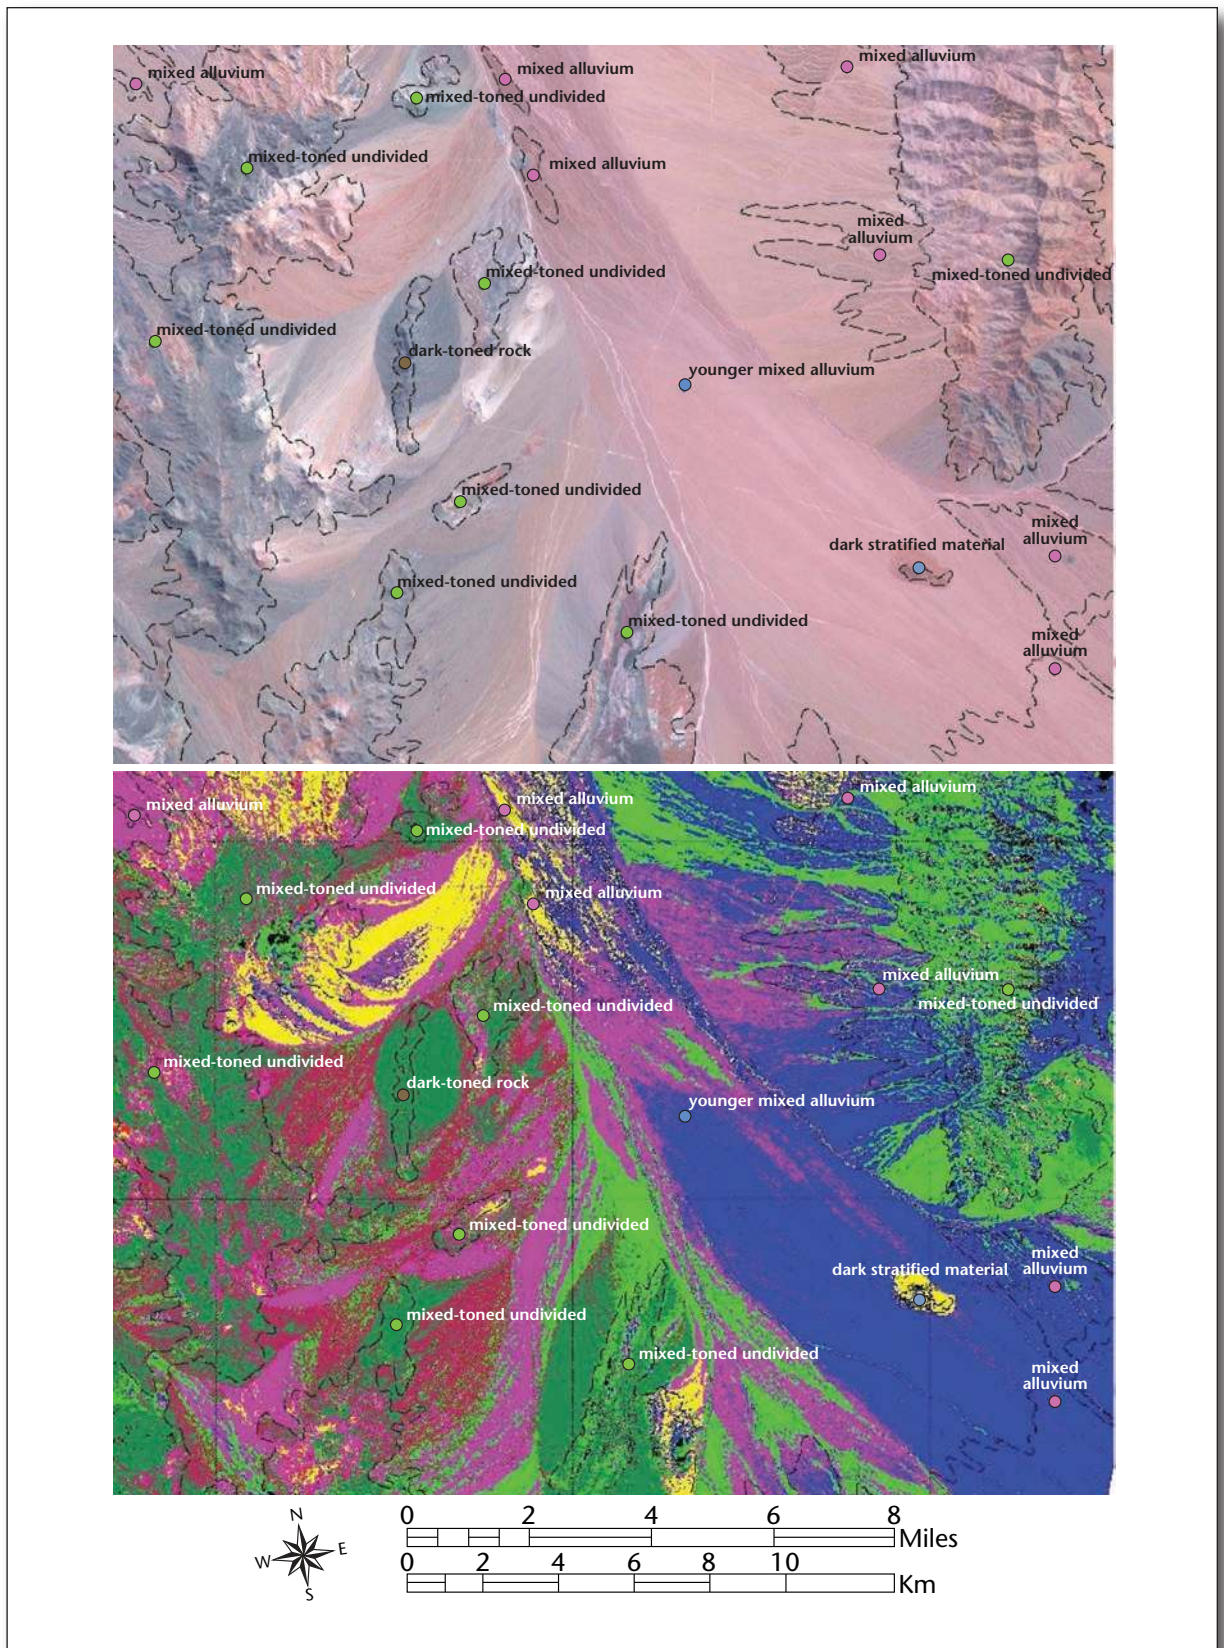
Figure 3. Examples of the Different Data Products.

Top: Landsat image used in traverse planning. Bottom: Geologic classification map derived from ASTER data.

Figure 4. Hazard Avoidance.

Rover Platform, Instruments, and Operations

In typical field work, rover operations follow a daily cycle in which a remote science team reviews the prior data, decides the next day's navigation waypoints and measurements, and then sends these commands to the rover over a satellite link. This is similar to the sporadic communications of a planetary mission. The rover then executes its commands over the course of the subsequent day. In the Atacama campaign, typical command cycles for Zoë cover 5–10 kilometers per day. Figure 2 shows the entire traverse path: red dots show locations for imaging and spectral data collection, while white paddles indicate sites of particular interest where more in-depth study is performed.

Scientists determine the waypoints for the next day using geologic and compositional maps produced from orbital remote sensing data. Here the ASTER instrument proves particularly useful: its images have a spatial resolution of 15 meters (visible) and 30 meters (SWIR), making them capable of resolving details such as isolated rock outcrops.

While the three visible and six SWIR bands are not sufficient to conclusively identify mineralogical composition, they do help discriminate the principal units of surface material and suggested representative sites to visit.

Figure 3 shows examples of the different data products: a Landsat image with three visible bands reveals terrain morphology and desirable outcrops, and a multiband ASTER image provides a rough classification of mineralogical units.

The rover itself is capable of driving more than 10 kilometers per day on challenging desert terrain (Wettergreen et al. 2008). On-board obstacle avoidance uses three-dimensional geometry from stereo imagery to identify hazards above the ground plane and plan local drive arcs that go around them (figure 4). Figure 5 shows the robot and the components used by its science autonomy system. A pair of forward-facing navigation cameras provide hazard avoidance capability through a local path planner. The vertical drill structure delivers subsurface soil to a microscopic imager and a Raman spectrometer inside the rover. Analyzing the drill samples takes an hour or more, so these are deployed judiciously at specific locations. However, we found that autono-

Figure 5. The Main Components of Zoë's Vis-NIR Spectrometer System.

A Raman spectrometer inside the rover measures pulverized samples from the subsurface drill.

my could play a role in improving the science data collected by the Vis-NIR spectrometer. The spectrometer is a modified Analytical Spectral Devices Fieldspec Pro that acquires radiance spectra from 0.4–2.5 micrometers at 0.001 micrometer resolution, housed in the rover body and connected by a fiber optic cable to a foreoptic telescope mounted on a pan-tilt mechanism. The foreoptic provides a 1 degree field of view, and can be directed at specific targets in the environment. Its field of regard spans a full 360 degrees azimuth and 90 degrees elevation. A colocated camera provides visual context to interpret the spectra.

Zoë's Vis-NIR reflectance data overlaps in wavelength with ASTER orbital images; it is a more spatially and spectrally refined version of the satellite data. By visiting distinctive terrain units of figure 3,

analysts can refine the remote view with detailed spectral information and specific mineral absorption features. In this manner the Vis-NIR data serves as both a validation of the orbital data and a means to better interpret the mineralogical constraints and context for biogeology studies. Each session of Vis-NIR acquisitions begins with the rover calibrating its instrument for temperature, solar geometry, and atmospheric conditions using a white reference target mounted on the rover deck (figure 5 inset).

Dividing the radiance from the target by the reference measurement produces reflectance data of the form shown in figure 6. These spectra were acquired at locations indicated in the adjacent panoramic camera subframe, from a distance of approximately 2 meters. The reflectance values represent the fraction of light reflected at each wavelength; more specific
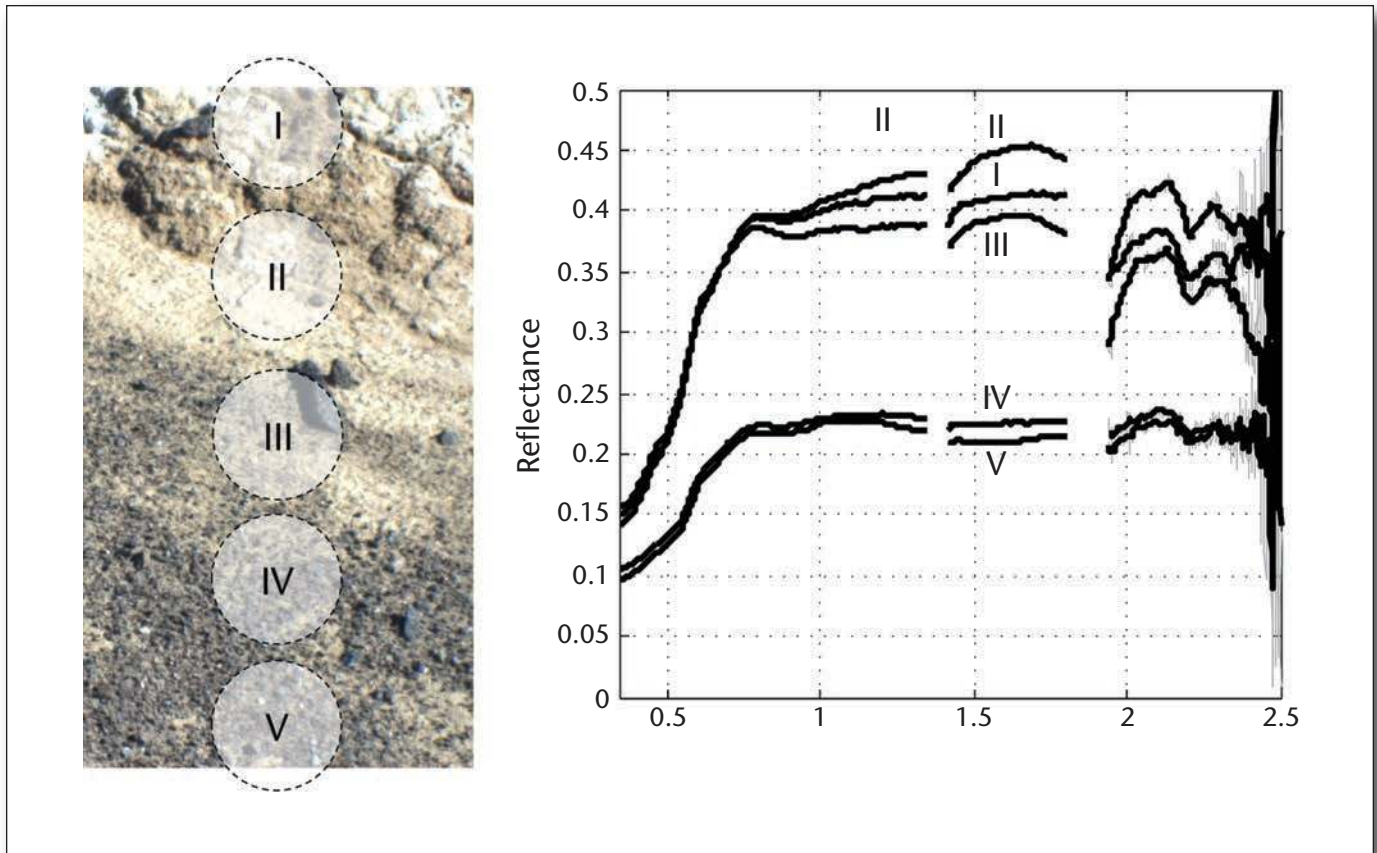
Figure 6. Panoramic Camera Subframe.

Dive spectrometer fields of view (PP24), and associated reflectance spectra.

formulations are possible (Schaepman-Strub et al. 2006), but we will use reflectance here in the ordinary Lambertian sense. This assumption should generally hold for the geologic materials of interest. Note that the light-colored sediment area in spectra I-III is associated with a higher average reflectance, as well as unique spectral features such as the dip near 2 micrometers. These spectra were smoothed using local linear regression, but some lingering noise spikes at longer wavelengths evidence the lower signal level in these spectral regions.

Science Autonomy Methods

Zoë's science autonomy system includes two basic capabilities that operate on mesoscale and macroscale features respectively. Smart targeting can identify science features in rover navigation imagery and use this information to point the Vis-NIR spectrometer. Adaptive path planning navigates on scales of tens or hundreds of meters, using satellite images to select waypoints with distinctive or novel spectra. We describe each of these techniques in turn.

Smart Targeting

Zoë began each autonomous target selection process by acquiring a navigation camera image. On-board image processing then analyzed the scene to find large contiguous regions of a desired terrain class. Typically these classes were rough surface features like rock outcrop or bright sediment patches with distinctive spectral signatures. Upon finding a feasible target, the rover recalibrated its Vis-NIR spectrometer, pointed at the feature, and collected a small 3 x 3 raster of spectra centered on the target of interest. For context, it also acquired a high-resolution color image of the scene.

The image analysis used a random forest pixel classification system described in previous work (Foil et al. 2013; Wagstaff et al. 2013) and adapted to the Atacama environment. This supervised classification method learns a mapping from local pixel intensities to the surface class of that pixel. The model is instantiated as an ensemble of decision trees trained in advance. At run time, the rover tested each pixel in the new image and averaged the classification of each tree in the ensemble. The end result was a classification map of the entire image, along with asso-

ciated class posterior probabilities. By subsampling each image by a factor of four prior to classification, processing time was less than a second on Zoë's onboard laptop-scale CPU.

After image classification, connected components analysis was used to identify contiguous targets. The rover then promoted the single largest target of the desired class for followup data collection. For each target, a center pixel was determined using the largest inscribed circle heuristic (Estlin et al. 2012) and transformed to a pan-tilt angle using the assumption of a planar terrain surface. Use of navigation camera stereo data would identify a true three-dimensional position and enable more sophisticated kinematic solutions. Here we relied on an approximate planar solution coupled with rastering to ensure that an inaccurate pointing would still capture the target in at least one of the Vis-NIR spectra.

Scientists developed several different ways to incorporate this capability into multiday rover operations. The first approach was a target check used in the middle of long traverses. This only deployed the spectrometer if a feature of interest was found at the check point. If there was no feature in the rover's field of view, it would carry on without spending the time to calibrate and deploy its spectrometer. In this fashion, Zoë could cover long distances without spending undue time on bare or uninteresting terrain. This strategy was also useful near the boundary of geologic contacts where the precise location was uncertain. A second strategy involved a paired panorama that acted as a supplement to a commanded Vis-NIR spectral raster. Here the rover committed all time resources in advance. It calibrated its spectrometer and acquired data in columns of five spectra spaced at 10 degree increments directly in front of the robot and to either side. This provided a representative sampling of the terrain comprising the rover's current ASTER pixel. It then augmented this dataset with a 3 x 3 raster centered on any target of interest. Together, these two products gave a better insight than either taken individually. They met the dual needs of having representative spectra as well as capturing distinctive (outlier) features.

Adaptive Path Planning

The science autonomy system also operates on larger scales of tens or hundreds of meters, where it analyzes satellite data to adjust its traverse path. We model the explored environment using a standard geographic or area mixing model where each measurement is a mixture of a small number of end-member materials. End members' spectra combine in proportion to their physical extent on the surface. Most scenes contain just a few end-member spectra, and any measurement $\mathbf{x}$ can be reconstructed with appropriate constituents and mixing fractions. For a scene with m end members we define the mixing fractions to be vectors $\phi \in \mathbb{R}^m$. More generally we can model a spectral image using a linear combination of library spectra given by a d x m matrix Y. This gives the relationship $x = Y\phi$.

In practice there is always residual error separating the reconstruction from the measurement. This is partly attributable to measurement noise, but unless the library is comprehensive there may also be incompleteness errors (for example, spectral features that are expressed in the observations but not present in the library). A library that reconstructs all spectra well can be said to have explained the scene, and provides insight into the mineral compositions in the remote sensing data. This intuition provides a figure of merit for an adaptive path-planning system to select future measurement locations. Zoë's planner selects locations, the measurements at which, when used to augment the collected library Y, provide the largest expected reduction in unmixing error. The planner aims to visit locations that are spectrally distinctive, collecting samples that fully explain the orbital image.

In detail, we begin with a space of candidate measurement locations $\mathcal{L}$. The robot collects a library of spectra by sampling at sites $B = \{b : b \in \mathcal{L}\}$. We define a stochastic measurement function, $\mathbf{y} = f(b) + \varepsilon$ with Gaussian-distributed measurement noise ε, that yields spectral measurements $Y = \{\mathbf{y}_i : \mathbf{y}_i \in \mathbb{R}^d, 1 \leq i \leq m\}$. Together the observations form a spectral library, a random m x d matrix written Y_B. Good measurements reduce the total reconstruction error for selected remote sensing observations given by $X = \{\mathbf{x}_i : \mathbf{x}_i \in \mathbb{R}^d, 1 \leq i \leq n\}$. We impose a resource cost $C(B)$ to represent limited time, power, and bandwidth; it must not exceed a total budget β. For simplicity we will initially ignore the cost of point-to-point travel.

We define a risk function as the expected reconstruction error incurred from unmixing the remote images with the library of spectra collected at the surface:

$$R(B) = E\left[\min_{\phi} \sum_{\mathbf{x} \in X} \left\| Y_B \phi - \mathbf{x} \right\|_2 \right]$$
$$\text{for } \phi \geq 0,\ C(B) \leq \beta \tag{1}$$

Here we are taking the expectation over the rover's observation matrix, which is a random variable. Computing this expectation is analytically challenging, so instead we solve the related problem:

$$\arg\min_B = \min_{\phi} \sum_{\mathbf{x} \in X} \left\| X_B \phi - \mathbf{x} \right\|_2$$
$$\text{for } \phi \geq 0,\ C(B) \leq \beta \tag{2}$$

The linear geographic mixing assumption allows us to infer the expectation $E[YB]$ since in situ spectra combine in proportion to their extent on the surface. Due to geographic mixing we can directly substitute the remotely observed spectra as the expected observations. We rewrite the objective function using remote measurements at sites B, written X_B:
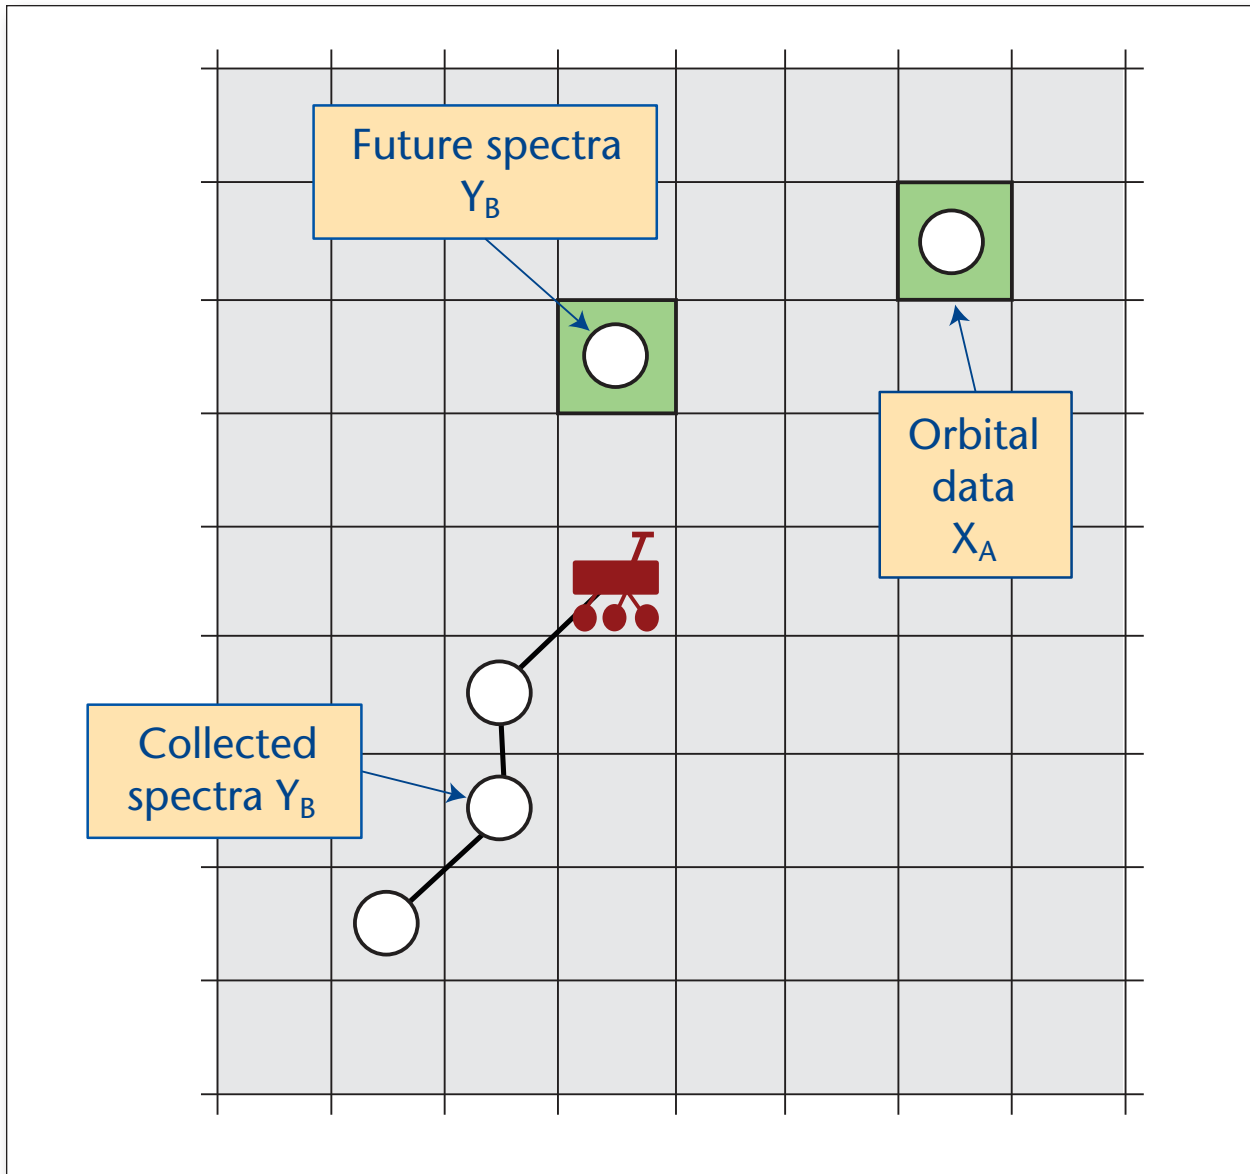
Figure 7. Formulation of Adaptive Path Planning.

$$R(B|A) = \min_{\phi} \sum_{\mathbf{x} \in X} \|X_B \phi - \mathbf{x}\|_2$$
$$\text{for } \phi \geq 0, C(B) \leq \beta \qquad (3)$$

This allows direct computation of the objective for any candidate set of measurement locations.

As the robot begins to collect spectra, some elements of $E[Y]$ become observed. The matrix Z_A represents the library of spectra collected at previous locations $A = \{a : a \in \mathcal{L}.\}$. These measurements are a realization of Y_A, and can be substituted into the expectation as the library of in situ spectra grows. Consequently the library used for unmixing consists of (1) the actual spectra collected at previous locations concatenated with (2) the expected spectra collected

at the future candidate locations. The objective is:

$$R(B|A) = \min_{\phi} \sum_{\mathbf{x} \in X} \|[Z_A X_B] \phi - \mathbf{x}\|_2$$
$$\text{for } \phi \geq 0, C(B) + C(A) \leq \beta \qquad (4)$$

To summarize, this final form reflects the key elements of Vis-NIR surface exploration: the overall goal of an accurate model using a handful of spectra, reflected in the squared error term; the physical behavior of geographic mixing, which appears as a positivity constraint; and the overall path-length budget β representing finite rover energy and time resources.

Figure 7 portrays the planning process. Here the robot has collected two spectra to form its library Z_A.

It calculates the expected risk of a candidate path using remote sensing data at locations X_B as a proxy for future measurements. In this manner, it can greedily (or nonmyopically) construct an optimal path. For our tests, the path planning was purely greedy; we added waypoints one by one, inserting them into the optimal location in the waypoint sequence and stopping when the total Euclidean path cost was exceeded.

During the Atacama field season we defined a budget defined in terms of path length, typically providing 1.5 times the straight-line distance to the goal. The time cost could be significant for longer traverses, particularly including the cost of the spectral measurements at intermediate waypoint. For this reason, most navigation actions were driving commands with adaptive navigation actions at particular regions of special interest. When it encounters a science waypoint in the plan, the science path-planning software finds the complete interpolating path that minimizes spectral reconstruction error of the nine-band ASTER image. It drives to the next intermediate waypoint along that path, collects a spectrum of the terrain, and replans the remainder using whatever path budget remains. That remainder becomes the next science plan, which is further refined in additional planning rounds as the rover progresses forward. In this fashion, the science planner can be fully stateless and react to new data encountered during the traverse.

Figure 8 shows the benefit of science-aware path planning in a simple simulation. We simulate a virtual rover traversing the famous Cuprite, Nevada, mining district, which is known for containing many distinctive spectral features of interest in highly localized outcrops (Swayze et al. 1992). Here we planned rover paths using orbital ASTER data, simulating 256 trials with random start and end points within the scene. We also simulated high-resolution in situ acquisitions using coregistered data from the Airborne Visible Near Infrared Spectrometer (AVIRIS) (Green et al. 1998).

The comparison considers four different strategies to fill the path-length budget: a random path, which bends the path toward randomly selected intermediate waypoints; a direct path, which begins with a straight line and then adds "wiggles" until the total length is reached; an unconstrained adaptive approach that minimizes equation 3 but without the positivity constraint; and an adaptive approach that enforces positivity of mixing fractions. We recomputed the reconstruction error for every trial by applying nonnegative least squares to the collected high-resolution spectra. Figure 8 shows each method's performance with a box indicating the median and quartile of the data and the whiskers indicating the extrema. Both adaptive methods significantly outperform the nonadaptive approaches, with the constrained adaptive method performing best of all. This performance boost happens because only the adaptive systems actively pursue the isolated outcrops with unique mineralogy.

On-board the real rover, a low-level control system is required to travel safely between these features of interest. Consequently, Zoë has on-board navigation software to turn high-level science waypoints, spaced on the order of tens or hundreds of meters, into low-level vehicle actions like drive arcs. It uses a software suite known as the Reliable Autonomous Surface Mobility (RASM) package (Wettergreen and Wagner 2012) capable of local hazard avoidance and path planning using a three-dimensional terrain representation. RASM extracts a cloud of three-dimensional points from the stereo cameras, orients these points relative to previously collected data, and builds a triangulated mesh. An A* search algorithm projects drive arcs across this mesh to compute the cost of local control actions. On longer scales, graph search identifies the best path to the next waypoint. RASM retains knowledge of the topology relating observation locations to their neighbors, permitting efficient loop closure and pose estimation over long distances.

Field Season Results

We engaged smart targeting during three days of rover operations. Table 1 shows the performance for typical targets during these traverses. Columns indicate the day (Sol) of operations; the action sequence number, with TC indicating a target check, PP a paired panorama, and AT a more specific planned data collection activity; the analysts' post hoc interpretation of the feature that was found; and two columns indicating whether the result was a reasonable science target and whether the pointing was accurate. Pointing accuracy was evaluated based on the context image collected with each spectrum, allowing it to be placed within the navigation camera image.

Overall, target selection performed reliably. The majority of targets were either rocks or patches of distinctive sediment. The only arguable failure was when the system identified a distant car that fell into the rover field of view. The pointing solution was slightly less reliable, since our groundplane assumption tended to break down at the periphery of the navigation camera image near the horizon.

Occasionally very distant targets would result in the rover aiming its spectrometer too far into the distance. Only one scene was totally featureless — an empty plain — and in this case the rover detected no targets and successfully abstained from spending any time resources.

Figure 9 shows several images from the real-time classification. Here the system was trained to recognize rocks and high albedo soil patches, and it successfully finds these features. In the center column, a red overlay represents the surface labeled as the belonging to the target class. The black rectangles

Figure 8. Adaptive Path-Planning Performance, in Simulation.

Sol	Action	SR	Target Found	Target Valid?	Pointing Accurate?	Notes
12	TC30	122	Rocks	OK	OK	
	TC31a	123	Foreground rock	OK	OK	
	TC31b	125	Foreground rock	OK	OK	
	TC32	128	Rock pile, sediment	OK	OK	
13	AT-13-09	172	Disturbed rocks and sediment	OK	OK	
	PP22	166	Distant rock patch	OK	BAD	1
	PP23	160	Distant rock patch	OK	OK	
	PP24	154	Foreground rocks	OK	OK	
	PP25	148	Foreground rocks	OK	OK	
	TC34	158	Foreground sediment patch / rocks	OK	OK	
	TC34-recon	132	None	OK	n/a	2
	TC41	170	Rocks	OK	BAD	3
	TC42	164	Distant rock patch	OK	BAD	4
14	AT-13-10	190	Car	BAD	BAD	5
	PP19	216	Foreground Rock	OK	OK	
	TC40	183	Rock patch	OK	OK	

Table 1. Target Detection Results from Playa Exploration Phase.

1: Aims too high for distant targets. 2: No target in scene. 3: Targeted feature not the largest rock. 4: Very distant feature. 5: Cars in frame.

show the field of view of the high-resolution follow-up image collected by the mast-mounted camera (right column). Each follow-up image is accompanied by a 3 x 3 raster of spectra. Even when the target selection was successful, we did not notice a significant difference between the on- and off-target spectra. This may be attributed to low signal to noise. Alternatively, these features may have spectral signatures that were very similar to the background substrate.

We deployed the adaptive navigation system successfully in two instances during the 2013 field campaign. Near the end of the field season the rover vis-

Figure 9. Examples of Targets Detected in Midtraverse, and Associated Followup Images.
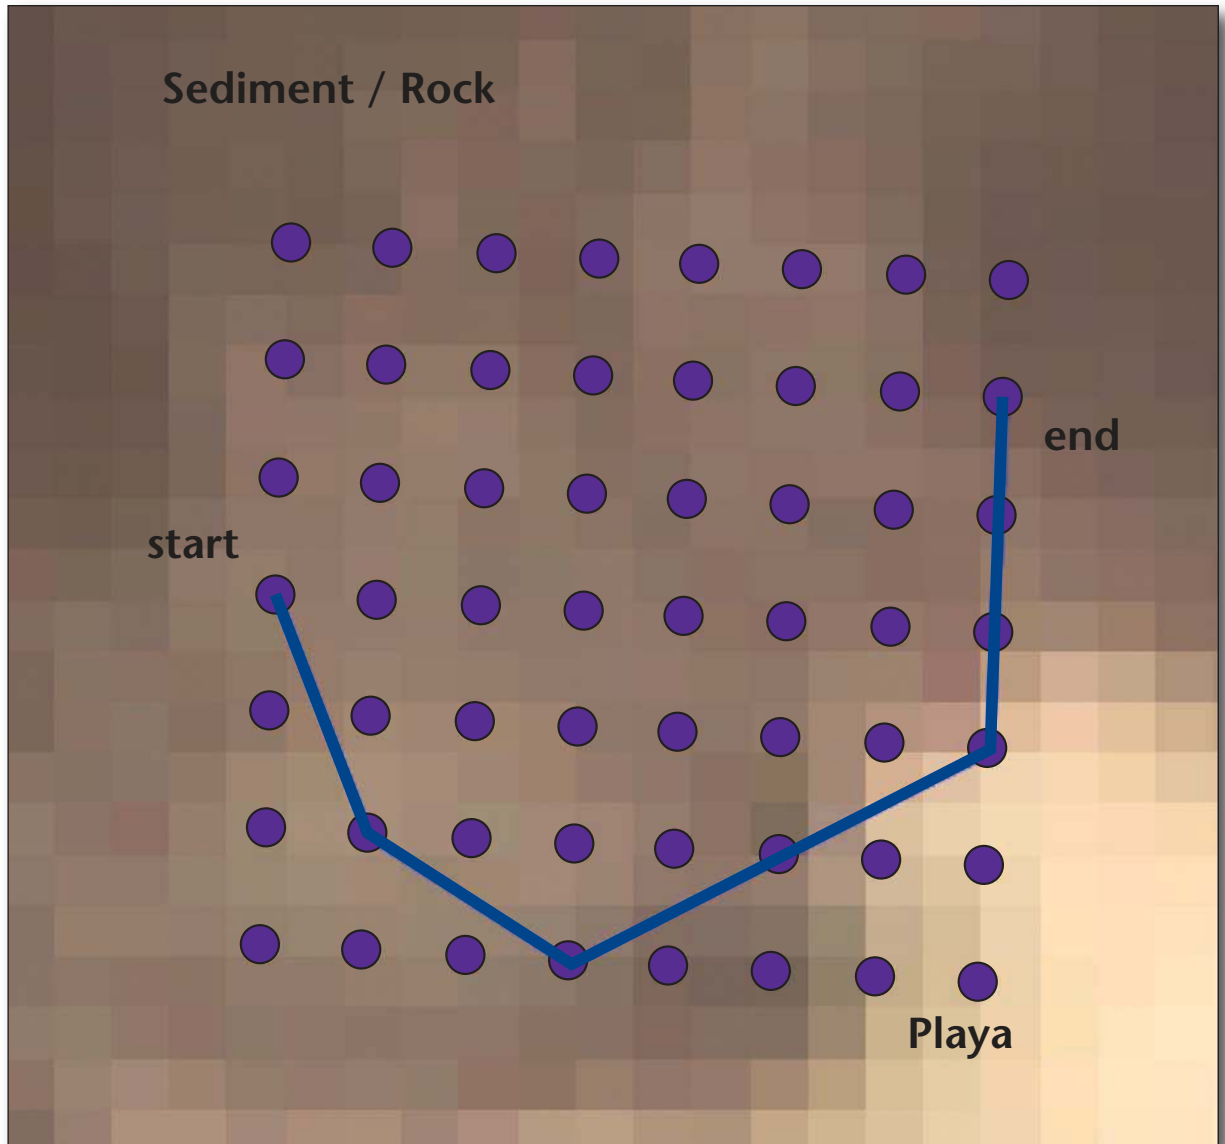

Figure 10. A Composite of Rover Images Showing the Playa Where Adaptive Navigation Was Evaluated.

ited a playa — a dry lakebed approximately 2 kilometers in length that was spectrally distinctive from the surrounding terrain (figure 10). Figure 11 shows a typical round of adaptive path planning near the playa edge. Here the playa is visible as a bright area in the lower right of the overhead satellite image. The large pixels represent 15-meter ASTER data. Here Zoë's planned path, in blue, diverts to sample the spectrally distinctive playa surface. The path changed only slightly in subsequent replanning as the rover visited each waypoint and incorporated the new spectra into its solution.
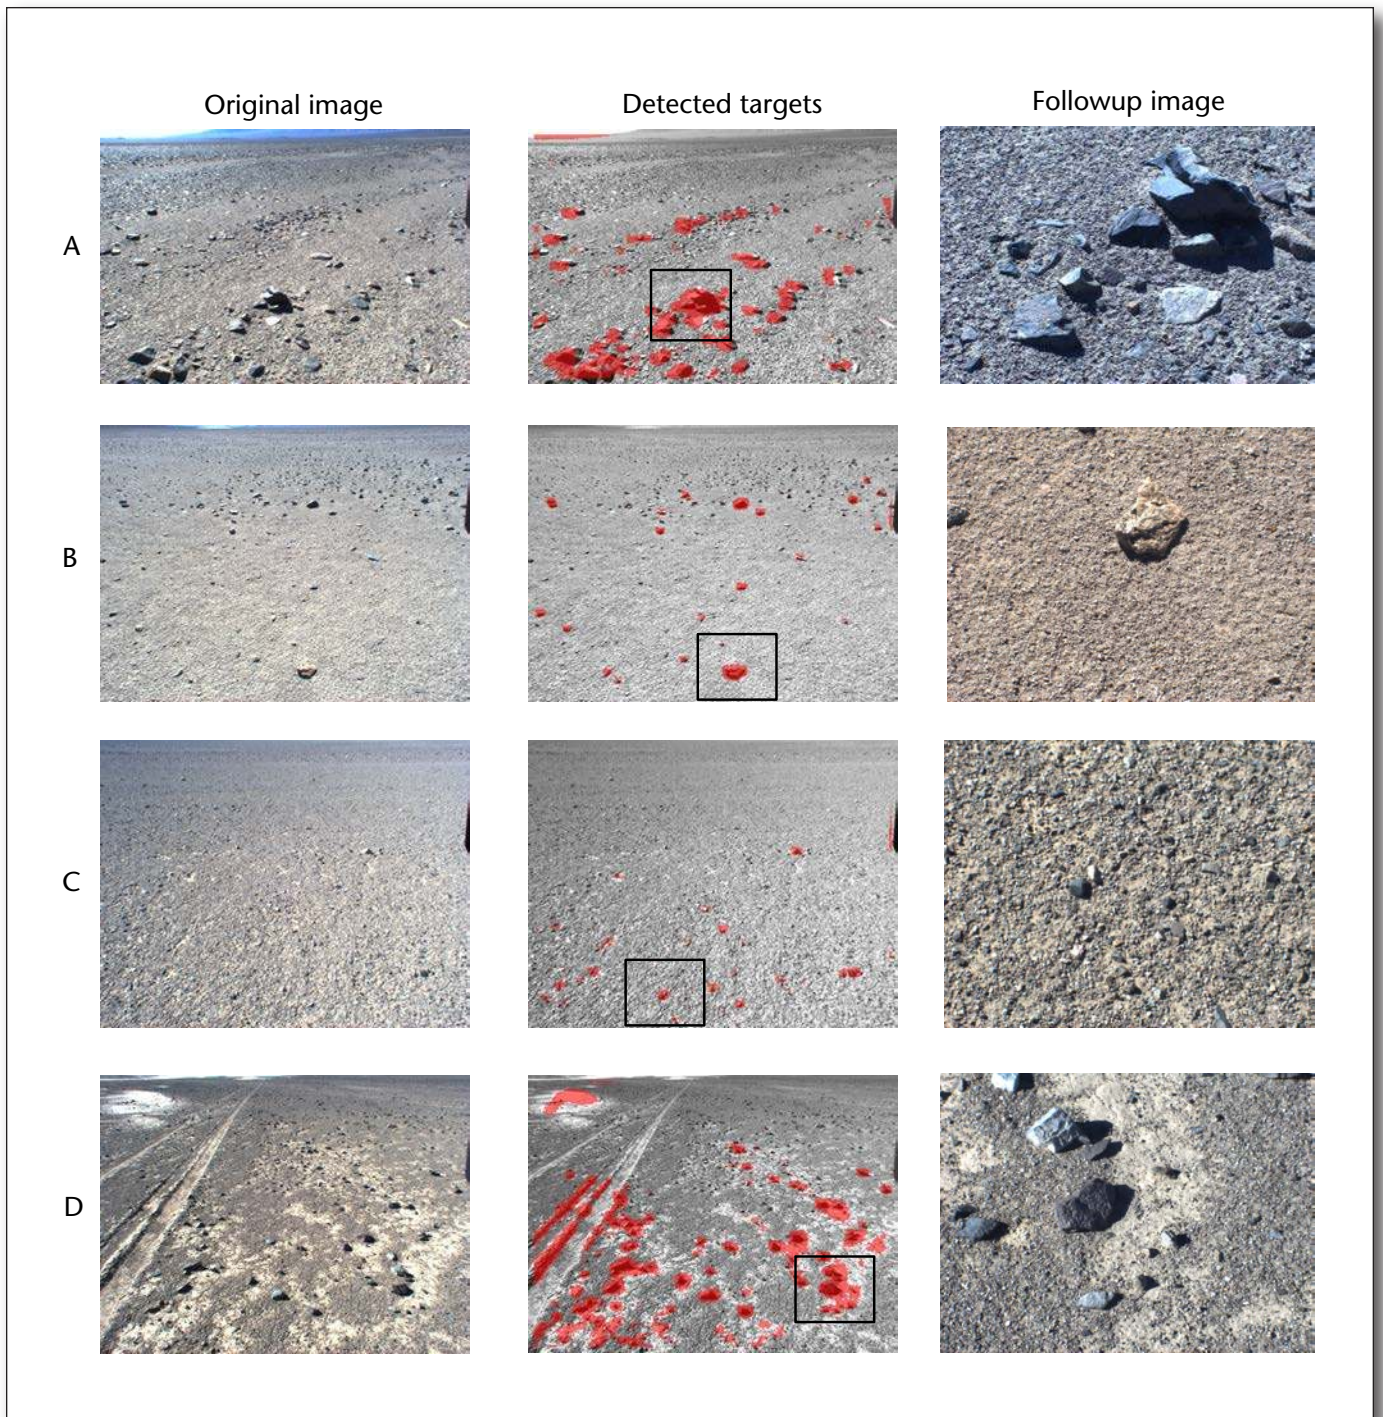
	Original image	Detected targets	Followup image

Figure 11. Demonstration of Adaptive Path Planning.

While the on-board planning gave an intuitive and reasonable answer, the actual rover paths were not as expected due to misregistration between orbital data products and the rover's on-board GPS estimate. Postanalysis of the data revealed the real position was offset by more than 100 meters from the intended location, so the actual rover path spent most of its time on the playa. In the future we will directly address these registration errors with the use of explicit ground control points (GCPs).

Conclusions

This work demonstrates novel techniques integrating adaptive autonomous science activities with pre-planned data collection. Zoë's system will continue to mature in the coming year.

Acknowledgments

References

Cabrol, N. A.; Wettergreen, D.; Warren-Rhodes, K.; Grin, E. A.; Moersch, J.; Diaz, G. C.; Cockell, C. S.; Coppin, P.; Demergasso, C.; Dohm, J. M.; Ernst, L.; Fisher, G.; Glasgow, J.; Hardgrove, C.; Hock, A. N.; Jonak, D.; Marinangeli, L.; Minkley, E.; Ori, G. G.; Piatek, J.; Pudenz, E.; Smith, T.; Stubbs, K.; Thomas, G.; Thompson, D.; Waggoner, A.; Wagner, M.; Weinstein, S.; and Wyatt, M. 2007. Life in the Atacama: Searching for Life With Rovers (Science Overview). *Journal of Geophysical Research: Biogeosciences* 112(G4): 28, 2007. dx.doi.org/10.1029/2006JG000298

Castaño, A.; Fukunaga, A.; Biesiadecki, J.; Neakrase, L.; Whelley, P.; Greeley, R.; Lemmon, M.; Castano, R.; and Chien, S. 2008. Automatic Detection of Dust Devils and Clouds on Mars. *Machine Vision and Applications* 19(5-6): 467–482. dx.doi.org/10.1007/s00138-007-0081-3

Chien, S.; Sherwood, R.; Tran, D.; Cichy, B.; Rabideau, G.; Castano, R.; Davies, A.; Mandl, D.; Frye, S.; Trout, B.; Shulman, S.; and Boyer, D. 2005. Using Autonomy Flight Software to Improve Science Return on Earth Observing One. *Journal of Aerospace Computing, Information, and Communication* 2(4): 196–216. dx.doi.org/10.2514/1.12923

Davies, A.; Chien, S.; Baker, V.; Doggett, T.; Dohm, J.; Greeley, R.; Ip, F.; Cichy, B.; Rabideau, G.; Tran, D.; Sherwood, R. 2006. Monitoring Active Volcanism With the Autonomous Sciencecraft Experiment on EO-1. *Remote Sensing of Environment* 101(4): 427–446. dx.doi.org/10.1016/j.rse.2005.08.007

Doggett, T.; Greeley, R.; Chien, S.; Castano, R.; Cichy, B.; Davies, A.; Rabideau, G.; Sherwood, R.; Tran, D.; Baker, V.; et al. 2006. Autonomous Detection of Cryospheric Change with Hyperion On-Board Earth Observing-1. *Remote Sensing of Environment* 101(4): 447–462. dx.doi.org/10.1016/j.rse.2005.11.014

Estlin, T.; Bornstein, B.; Gaines, D.; Anderson, R. C.; Thompson, D. R.; Burl, M.; Castano, R.; and Judd, M. 2012. Aegis Automated Targeting for Mer Opportunity Rover. *ACM Transactions on Intelligent Systems Technology* (TIST) 3(3): Article 50. dx.doi.org/10.1145/2168752.2168764.

Foil, G.; Thompson, D. R.; Abbey, W.; and Wettergreen, D. S. 2013. Probabilistic Surface Classification for Rover Instrument Targeting. In *Proceedings of the 2013 IEEE International Conference on Intelligent Robots and Systems* (IROS), 775–782. Piscataway, NJ: Institute of Electrical and Electronics Engineers. dx.doi.org/10.1109/IROS.2013.6696439

Green, R. O.; Eastwood, M. L.; Sarture, C. M.; Chrien, T. G.; Aronsson, M.; Chippendale, B. J.; Faust, J. A.; Pavri, B. E.; Chovit, C. J.; Solis, M.; Olah, M. R.; Williams, O. 1998. Imaging Spectroscopy and the Airborne Visible/Infrared Imaging Spectrometer (Aviris). *Remote Sensing of Environment* 65(3): 227–248. dx.doi.org/10.1016/S0034-4257(98)00064-9

Ip, F.; Dohm, J.; Baker, V.; Doggett, T.; Davies, A.; Castano, R.; Chien, S.; Cichy, B.; Greeley, R.; Sherwood, R.; Tran, D.; Rabideau, G. 2006. Flood Detection and Monitoring with the Autonomous Sciencecraft Experiment Onboard EO-1. *Remote Sensing of Environment* 101(4): 463–481. dx.doi.org/10.1016/j.rse.2005.12.018

Mastrodemos, N.; Kubitschek, D. G.; and Synnott, S. P. 2005. Autonomous Navigation for the Deep Impact Mission Encounter with Comet Tempel 1. *Space Science Reviews* 117(1–2): 95–121. dx.doi.org/10.1007/s11214-005-3394-4

Schaepman-Strub, G.; Schaepman, M.; Painter, T.; Dangel, S.; and Martonchik, J. 2006. Reflectance Quantities in Optical Remote Sensing: Definitions and Case Studies. *Remote Sensing of Environment* 103(1): 27–42. dx.doi.org/10.1016/j.rse.2006.03.002

Swayze, G.; Clark, R.; Sutley, S.; and Gallagher, A. 1992. *Ground-Truthing Aviris Mineral Mapping at Cuprite, Nevada. Summaries of the Third Annual JPL Airborne Geosciences Workshop, Volume 1: AVIRIS Workshop*. JPL Publication 92-14, 47–49. Pasadena, CA: Jet Propulsion Laboratory, California Institute of Technology.

Thompson, D.; Bornstein, B.; Chien, S.; Schaffer, S.; Tran, D.; Bue, B.; Castano, R.; Gleeson, D.; and Noell, A. 2013. Autonomous Spectral Discovery and Mapping Onboard the EO-1 Spacecraft. *IEEE Transactions on Geoscience and Remote Sensing* 51(6): 3567–3579. dx.doi.org/10.1109/TGRS.2012.2226040

Wagstaff, K.; Thompson, D. R.; Abbey, W.; Allwood, A.; Bekker, D. L.; Cabrol, N. A.; Fuchs, T.; and Ortega, K. 2013. Smart, Texture-Sensitive Instrument Classification for In Situ Rock and Layer Analysis. *Geophysical Research Letters* 40(16): 4188–4193. dx.doi.org/10.1002/grl.50817

Wettergreen, D., and Wagner, M. 2012. Developing a Framework for Reliable Autonomous Surface Mobility. Paper presented at the 12th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS-12), Turin, Italy, 4–6 September.

Wettergreen, D.; Wagner, M. D.; Jonak, D.; Baskaran, V.; Deans, M.; Heys, S.; Pane, D.; Smith, T.; Teza, J.; Thompson, D. R.; Tompkins, P.; and Williams, C. 2008. Long-Distance Autonomous Survey and Mapping in the Robotic Investigation of Life in the Atacama Desert. Paper presented at the 9th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS-08), Los Angeles, CA USA, 26–29 February.

Woods, M.; Shaw, A.; Barnes, D.; Price, D.; Long, D.; and Pullan, D. 2009. Autonomous Science for an Exomars Roverlike Mission. *Journal of Field Robotics* 26(4): 358–390. dx.doi.org/10.1002/rob.20289

David Wettergreen is a research professor at the Robotics Institute at Carnegie Mellon University. His research focuses on robotic exploration underwater, on the surface, and in air and space; and in the necessary ingredients of perception, planning, learning and control for robot autonomy. His work spans conceptual and algorithm design through field experimentation and results in mobile robots that explore the difficult, dangerous, and usually dirty places on Earth, in the service of scientific and technical investigations. Much of this work is relevant to ongoing and future space activities.

Greydon Foil is a Ph.D. student at the Robotics Institute at Carnegie Mellon University. He is interested in large-scale mapping and planning for autonomous rovers, with current work involving physical process modeling for better terrain understanding.

Michael Furlong is a researcher with interests in theoretical neuroscience, decision making, and robotics.

David R. Thompson is a researcher at the Jet Propulsion Laboratory, California Institute of Technology. His work involves the application of computer vision, machine learning, and spectroscopy to robotic Earth and planetary science. His algorithms have run on autonomous robots and sensors fielded in North America, South America, the Atlantic Ocean, airborne demonstrations, low Earth orbit, and the surface of Mars.