AUTONOMOUS ROVER DETECTION AND RESPONSE APPLIED TO THE SEARCH FOR LIFE VIA CHLOROPHYLL FLUORESCENCE IN THE ATACAMA DESERT. T. Smith, D. R. Thompson, S. Weinstein, D. Wettergreen, [trey,drt,sw42,dw0s]@andrew.cmu.edu. The Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh PA 15217, USA..

Introduction: As planetary rovers travel longer distances in each command cycle, science conducted during traverse becomes a significant part of the overall mission. During traverse, the rover may pass an interesting feature never detected before (because it was not close enough) and never visited again (because continued progress was deemed more important than returning for further study). In this situation, we would like the rover to collect some information before it passes on. Like a good field scientist, it needs to detect interesting features and autonomously respond with followup observations [1,3].

This paper describes autonomous detection and response capabilities used to enhance a robotic search for life in the Atacama Desert of Chile, part of the NASA ASTEP Limits of Life in the Atacama campaign [4,5]. In the coastal desert region studied (site D, rh 10-80%), it was expected that sparse patches of lichens and cyanobacteria would be exposed on the surface. To study these organisms, our rover carried a fluorescence imager (FI), which could detect chlorophyll via its inherent fluorescence. It could detect other organic chemicals, such as DNA and proteins, after applying fluorescent marker dyes. The current design of the FI was optimized to study terrestrial life, but similar instruments could be deployed on Mars or Europa, possibly focusing on a different set of biomarkers.

Application of FI dyes was time-consuming and used up a limited supply of dye, so it was important to ensure each dye sample contained interesting data. Thus, we focused the dyes on samples where chlorophyll was autonomously detected.

Autonomous chlorophyll detection and followup dye application was field tested in the Atacama. Samples autonomously chosen for dye followup were significantly more likely to contain photosynthetic organisms than samples chosen at random. Several of the autonomous followup images showed positive dye fluorescence. The new autonomy capabilities were successfully integrated into the operations scheme and saw continuing use by the remote team operating the rover.

Fluorescence imager: The FI is a down-pointing camera mounted on the bottom of the Zoë rover (fig. 1). It has 10 cm field of view and resolution 180 μm. During autonomous response experiments, the sampling location under the FI was chosen by stopping the rover at fixed distances along its traverse, and the camera was deployed and auto-focused using z-axis motion. The sample under study was not moved.

The FI could be used to detect either the reflectance or fluorescence of a sample in various channels. A xenon flashlamp provided illumination. Six optical interference filters could be automatically switched into the excitation path between the flashlamp and sample, and another ten filters could be switched into the detection path between the sample and CCD.

The FI captured reflectance under a combination of sunlight and flashlamp illumination with no excitation filter. Separate images with red, green, and blue emission filters were combined to form a visual color image. In fluorescence mode it captured a black and white intensity image, with excitation and emission channel pair selected to respond to the fluorescence of the chemical species under study, either chlorophyll or an artificial dye whose fluorescence greatly increased when it was bound with a biomarker. Different marker dyes responded to DNA, proteins, lipids, and carbohydrates. The FI rejected ambient sunlight during fluorescence imaging by limiting CCD exposures to 20 μs, synchronized with the flash. An automatic sprayer could spray the sample either with water (which enhanced chlorophyll fluorescence under dry conditions) or with a solution that contained all four marker dyes along with acid and detergent to aid dye penetration.

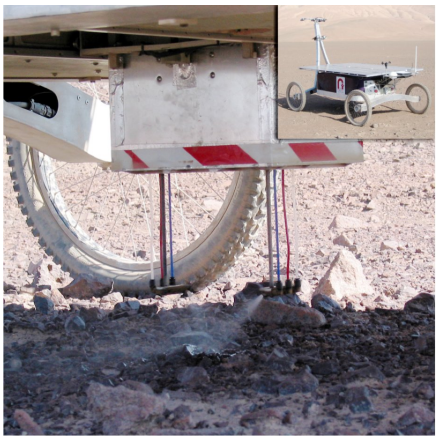

Figure 1: The FI deployed and spraying underneath the robot Zoë in the Atacama Desert of Chile with an overlay of Zoë in the desert.

Experimental procedure: Autonomous followups were enabled during 180 m traverses with periodic stops for FI samples (fig. 2). At each endpoint of the traverse, the rover stopped and took a *full sample*. At 30 m intervals within the transect, the rover stopped and took a *periodic sample*. The rover executed each traverse autonomously within a single command cycle (including both driving and FI samples).

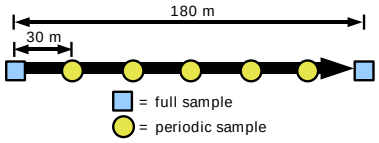

Figure 2: Standard sampling traverse.

The protocol for each sample had two phases: (1) The FI sprayed water and captured several images, including a chlorophyll image used to determine if followup was warranted. (2) It sprayed the dye mixture and captured several images, including dye fluorescence images. Execution times for the phases were approximately 8 and 15 minutes, respectively. For a full sample, both phases were always executed. For a periodic sample, the FI always executed phase 1, but executed phase

2 only if chlorophyll was autonomously detected (fig. 3). In addition, there was a followup quota: over the course of a traverse, only the first three periodic samples with positive chlorophyll detections could trigger a followup. This quota reduced time uncertainty during mission planning.

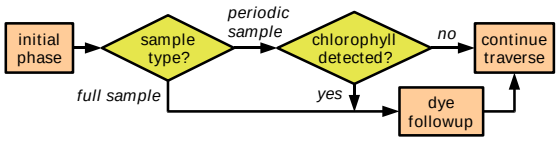

Figure 3: Sample protocol flowchart.

Chlorophyll detection: Autonomous chlorophyll detection, used only for periodic samples, relied on a single image of chlorophyll fluorescence intensity (excitation 450 nm, emission 740 nm) captured after the FI sprayed water on the sample. This was called the *trigger image* (fig. 4).

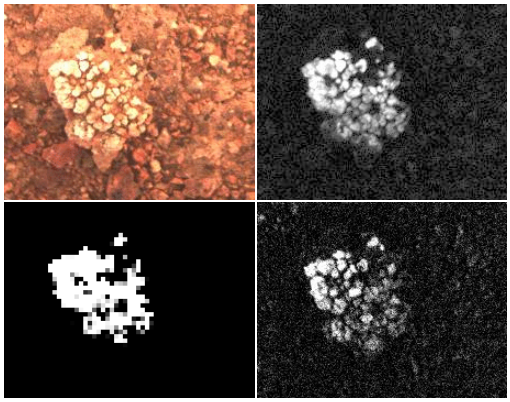

Figure 4: (top left) Portion of FI visual color image containing a lichen. (top right) Chlorophyll trigger image. (bottom left) An intermediate step of image processing; the brightness in each cell represents the estimated probability that it contains chlorophyll. (bottom right) After autonomous followup, the FI detected fluorescence from the DNA marker dye.

The detection algorithm reported the probability that chlorophyll was present anywhere in the trigger image, triggering an autonomous dye followup if the probability was 50% or higher. The algorithm reported a high probability if there were any bright patches in the image. In detail: First, it split the image into 4×4 pixel cells and calculated an average intensity over each cell. This smoothing eliminated false detections from single-pixel shot noise. Second, it converted average intensity for each cell to a probability that the cell contained chlorophyll using a logistic or "fuzzy threshold" function. Finally, it calculated the probability of chlorophyll being present anywhere in the image by combining the probabilities from individual cells using naïve Bayes [2].

The logistic function used to convert cell intensity to probability of containing chlorophyll had two tunable parameters. In order to set these parameters, two trigger images containing lichens were hand-labeled; each 4×4 cell was labeled as either containing chlorophyll or not, based on morphological cues from both the trigger image and an associated visual color image. The parameters were set using logistic regression so as to maximize the likelihood of the given labeling.

The prior probability of any cell containing chlorophyll also needed to be specified (equivalently, the proportion of the ground surface expected to be covered by photosynthetic organisms). Based on a set of training images, the prior value was informally hand-tuned so that most images would fall on the correct side of the followup threshold (negative images below 50%, positive images above 50%). The resulting value of 0.005 was used for all of the reported results.

Results: The effectiveness of the autonomous followup system was evaluated over 24 periodic samples collected during five traverses. Since the standard traverse length was 180 m, and periodic samples were taken at 30 m intervals, nominally each traverse should have included five periodic samples. In practice, the number varied because the specified endpoints were not exactly 180 m apart. The last traverse was cut short due to nightfall after only two periodic samples. In two of the traverses, the rover filled the followup quota before the last periodic sample was taken; periodic samples after that were not included in the analysis because the quota prevented the rover from triggering further followups.

Each sample image set was analyzed by a remote team that included field biologists and fluorescence experts. Using both the visual color image and the trigger image, they labeled the samples as positive (does contain photosynthetic organisms) or negative (does not). The scientist labels were compared to the autonomous followup response. The results agreed for 19 of the 24 samples. 8 of the 24 samples were positive; 7 of the 11 samples chosen for autonomous followup were positive. Thus, samples chosen for followup were 91% more likely to contain photosynthetic organisms than samples chosen at random (significance level < 0.01 using one-tailed Fisher's exact test). 9 of the 11 samples chosen for followup showed fluorescence with at least one of the marker dyes.

The system had a number of correctable failure modes. Zoë's solar panels normally shaded The FI field of view, but in early morning or late evening it was exposed to direct sunlight, overwhelming the fluorescence signal. High winds at the study site sometimes blew the mist from the sprayer out of the FI field of view before it reached the ground.

Conclusions: Autonomous detection and followup were shown to be feasible in the context of real science data and integrated rover operations. We plan to extend these results broadly to other instruments and more interesting types of autonomous followup, such as handoff from navigation cameras to point spectrometers and contact sensors.

Acknowledgments: Thanks to Nathalie Cabrol, Joseph Flowers, Stuart Heys, Dominic Jonak, James Teza, Michael Wagner, Kimberly Warren-Rhodes, and Christopher Williams. This work was supported by NASA under grants NNG0-4GB66G, NAG5-12890, and NAG5-12833.

References: [1] Castaño R. et al. (2005) *Proc. IEEE Aerospace.* [2] Mitchell T. (1997) *Machine Learning*, McGraw-Hill. [3] Smith T. et al. (2005) *Proc. IEEE Aerospace.* [4] Waggoner et al. (2005) *LPS.* [5] Wettergreen D. et al. (2005) *Proc. iSAIRAS.*