AUTONOMOUS DETECTION OF NOVEL BIOLOGIC AND GEOLOGIC FEATURES IN ATACAMA DESERT ROVER IMAGERY.

D. R. Thompson[1], T. Smith[1], and D. Wettergreen[1] {drt,trey,dsw}@ri.cmu.edu
[1]The Robotics Institute, Carnegie Mellon University, Pittsburgh PA 15213. USA.

Introduction: Planetary scientists will soon benefit from a new generation of exploration rovers that can travel long distances in a single communications cycle. Much of this terrain will never be seen by humans so these rovers must autonomously identify interesting features for data collection [1]. They must also be able to prioritize the most important data for transmission to Earth. When a tightly-controlled teleoperation loop is infeasible autonomous data understanding lets scientists explore more terrain while optimizing the quality of the returned data [2,3].

Here we investigate context-sensitive models for onboard data analysis. These "science maps" consider not only the raw data but also the sampling location. They improve onboard data understanding by revealing environmental trends and boundaries that can inform autonomous sampling and data return decisions. They also help identify novel features by highlighting anomalies that are unexpected in context of the local environment. In this work, tests with navigation imagery suggest that context-sensitive data analysis using a Hidden Markov Model [4] offers performance benefits for novelty detection during rover traverse.

Technical Approach: We focus on an image sequence from an autonomous rover traverse in the Atacama Desert (a Mars-analogue desert in Chile) [5]. The rover platform for these tests was Zoë, an exploration robot developed at Carnegie Mellon [1]. Zoë is capable of navigating autonomously at up to 1m/s for multiple kilometers. We chose a traverse of 1888 navigation images from remote science operations during the NASA ASTEP "Life in the Atacama" expedition. This drive began on a rocky hillside and descended into a basin covered in finer material.

Image Processing The 320x240 pixel color images were captured at two-meter intervals from cameras mounted on the rover mast. They provide a 60° field of view of the terrain in front of the rover. We split the foreground half of each image into a grid of 10-pixel cells; each cell constitutes an independent "sample" of the imaging site (Figure 1). The system computes numerical attributes for the color and texture of each cell so that an image provides 384 distinct samples in color-texture space. We use the cell's fractal dimension [6] as an efficient texture measure, and color normalize images using the "greyworld" strategy [7].

Model Parameters The sequential structure of rover traverse imagery makes it amenable to description by a Hidden Markov Model (HMM). The HMM uses "hidden states" to estimate the unobserved biologic or geologic characteristics of the rover's environment. This permits context-sensitive novelty detection; the rover can account for the expected local environment when evaluating an image's likelihood. HMMs can also identify environmental boundaries using techniques like the Viterbi algorithm [10].

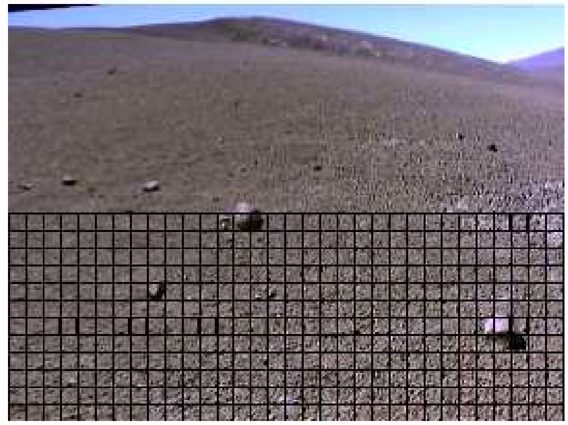

Figure 1: Navigation image. The algorithm extracts color and texture data from foreground grid squares.

A graphical representation of the HMM appears in Figure 2. The model quantifies relationships between neighboring states (transition probabilities) and the likelihood of each state generating a particular observed grid cell (emission probabilities). Its chain structure reflects the persistence of environmental conditions; the current state influences expectations for the next image's content. Our implementation utilizes a tied continuous- density Gaussian mixture model [8]. The Baum-Welsh algorithm [9] assigns model parameters which can then be used to estimate the hidden state for any image. We find parameters using the en-

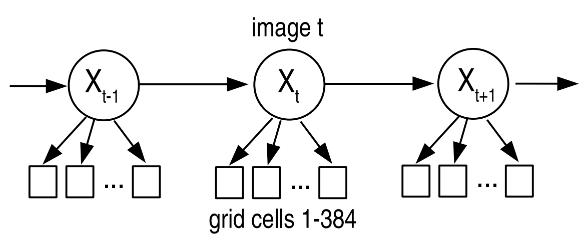

Figure 2: A graphical depiction of the HMM model. Arrows represent correlations between variables; at each time step t a hidden state X_t probabilistically generates a set of observed grid cells.

tire dataset but one could estimate models from partial data for adaptive sampling decisions during traverse.

Novelty Detection: With the complete data model one can find the likelihood of any grid cell in an image. Figure 3 identifies grid cells corresponding to the most unlikely image cells; these suggest cites for additional autonomous data collection.

More generally one can identify novel *images* for prioritized transmission to Earth. In this work we compare different data models' performance in detecting novel images from the traverse dataset. Absolute novelty detection performance is inherently subjective but we can still perform a relative comparison by choosing a set of rare "target" image features that are unambiguous to identify. We use plants (appearing in 132 images) rocks with an axis longer than 30 pixels (31 images), and the rover shadow (28 images). If the rover prioritizes images in order of decreasing novelty, the number of targets which are chosen gives some indication of how a model's likelihood estimates align with human notions of novel image content.

We used random model initializations to generate 20 trials simulating two different novelty detection strategies: a two-state HMM with three mixture components and a stateless three-component mixture model. The stateless case is not context-sensitive; it ignores each image's sequence position and treats all data products as independent.

Figure 4 shows the results of the experiment with boxes illustrating the middle data quartile and notches the 95% confidence interval for the median. A random data return policy provides a performance baseline. Both novelty detection methods exhibit good selectivity for the chosen features. The context-sensitive HMMs provide the best overall performance for this

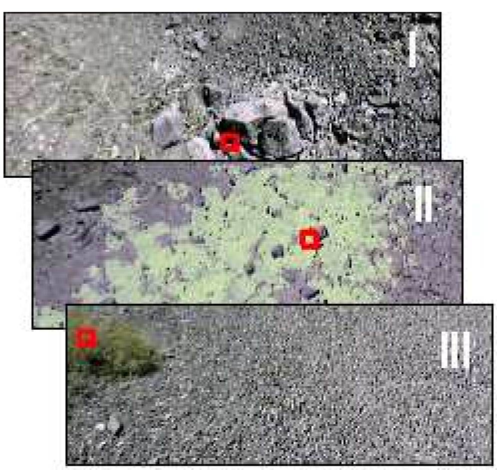

Figure 3: Examples of detected novel images: (I) large rock (II) high-albedo patch (III) plant life. Red squares indicate the most novel cell of each image.

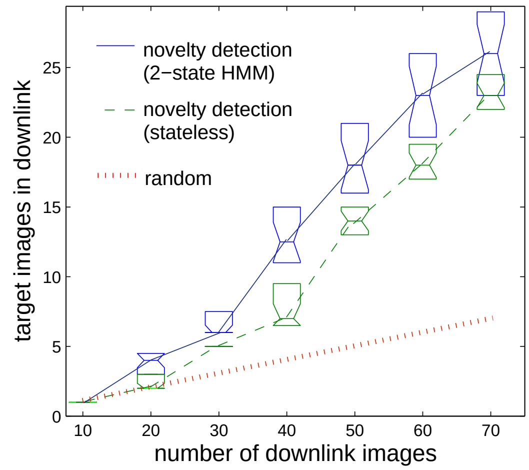

Figure 4: Performance graph showing selectivity for novel image features from the traverse.

dataset. Figure 3(II) shows one example of a novel image favored by the HMM. Individual cells from this high-albedo patch are similar to salt deposits visited earlier, but they are novel in context of the local environment.

Conclusion: HMMs have several characteristics that recommend them for autonomous science tasks dealing with sequential data. HMMs are theoretically well understood and computationally efficient. Preliminary tests suggest they may offer improved novelty detection performance.

The simple image analysis in this work contrasts with the complex pattern recognition we have favored in previous autonomous geology experiments. For planetary rover applications demanding computational efficiency, environment models may offer a complimentary path to improve science autonomy.

Acknowledgments: Thanks to Dominic Jonak, Joseph Flowers, Mike Wagner, Chris Williams, James Teza, and Stuart Heys. This research was supported by NASA grants NAG5-12890 and NNG0-4GB66G.

References: [1] Castaño R. L. et al. (2003) *Intl. Conf. Machine Learning.* [2] Castaño R. L. et al. (2003) *IEEE Aerospace*. [3] Wagstaff K. L. et al. (2004) *LPSC XXXV.* [4] Rabiner L. R. (1989) *Proc. IEEE 77:2, 257—286.* [5] Wettergreen D. et al. (2005) *ISAIRAS XIII.* [6] Chadhuri B. B. and Sarkar N. (1995) *IEEE Trans. Pattern Anal. and Machine Intel., 2:1, 72—77.* [7] Buchsbaum G. (1980). *J. Franklin inst., 310:1, 1—26.* [8] Bellegarda J. R. and Nahamoo D. (1990) *IEEE Trans. Accoustics, Speech and Signal Processing, 38:12, 1033—2045.* [9] Baum L. E. (1972), *Inequalities 3, 1—8.* [10] Forney G. D. *(1973) Proc. IEEE 61, 268—278.*