MARS ROVER EXPLORATION COMBINING REMOTE AND IN SITU MEASUREMENTS FOR WIDE-AREA MAPPING

Virtual Conference 19-23 October 2020

Alberto Candela[1], Kevin Edelson[1], David Wettergreen[1]

[1]The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA
E-mail: albertoc@andrew.cmu.edu, dsw@cmu.edu, kedelson@andrew.cmu.edu

ABSTRACT

We present an approach to autonomous rover exploration that enables higher science productivity. We first describe a machine learning model for wide-area mineral mapping that extrapolates signatures from just a few rover measurements. We use spectroscopic data because it is diagnostic of mineral composition. Exploration productivity is improved by incorporating notions from information theory and non-myopic path planning. We recently demonstrated the success of this approach in a field experiment in which our autonomous rover Zoë mapped a well-studied region of geological interest in Nevada. In this work, we apply and extend our methodology to actual Mars data and show performance in a Mars simulation study.

1 INTRODUCTION

Mars rovers have accomplished impressive science and exploration objectives. However, it is difficult to maintain high levels of productivity [1]. One of the reasons is the heavy reliance on interaction between the rovers and ground operators, requiring substantial effort when planning and validating commands that are sent to the robots. Furthermore, communication opportunities and bandwidth are limited, drawing out the time needed to accomplish mission goals.

The aim of this work is to improve mission productivity. First, by allowing rovers to efficiently map large areas through the extrapolation of scientific features of interest from just a few *in situ* measurements. Second, by enabling robotic explorers to take more adaptive actions based on real-time information.

In regard to efficient wide-area mapping, we use a machine learning model that allows a rover to map mineralogy by combining remote and *in situ* measurements. Spacecraft such as the Mars Reconnaissance Orbiter (MRO) and Mars Express collect remote data that has been crucial in the large-scale understanding of Mars, but often suffer from low resolution. Mars rovers such as Opportunity and Curiosity collect high-quality *in situ* data, but only at a few locations. Consequently, the objective of our model is to extrapolate rover measurements from *in situ* data in order to produce rich maps covering large areas. We focus on spectroscopic data since it is diagnostic in the study of surface composition and mineralogy on Mars; and consequently, it permits the assessment of geology, habitability, and potential biosignature presence [2, 3].

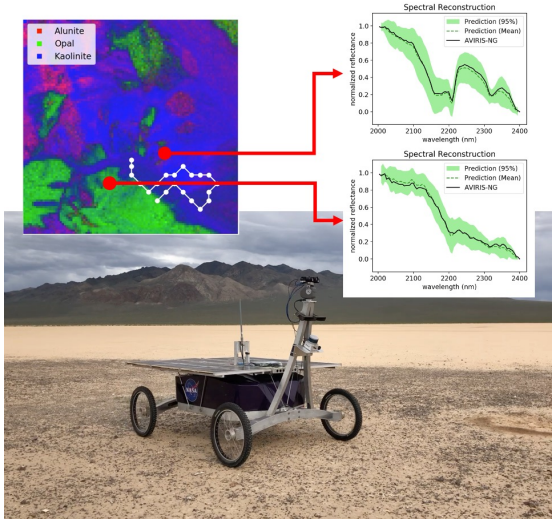

Figure 1: The autonomous rover Zoë performing a geological survey at Cuprite, Nevada. The robot collects samples (white path) with an onboard spectrometer and extrapolates spectral signatures to unsampled locations to then generate a mineral map.

This work also derives an exploration strategy for selecting samples that improve science productivity. We first use well-established principles from information theory, such as Shannon entropy [4], that allow the robot to identify the most meaningful science measurements [5, 6]. We then formulate rover exploration as an informative path planning problem where the goal is to minimize entropy, and which can be solved with a variety of existing algorithms [7, 8, 9].

The success of this approach was recently demonstrated [10]: the autonomous rover Zoë surveyed a well-studied region in Nevada (Figure 1). In this work, we extend and apply our methodology to actual Mars data and show performance in a simulation study that involves mineralogical investigations at two locations of interest: Jezero Crater and Nili Fossae.

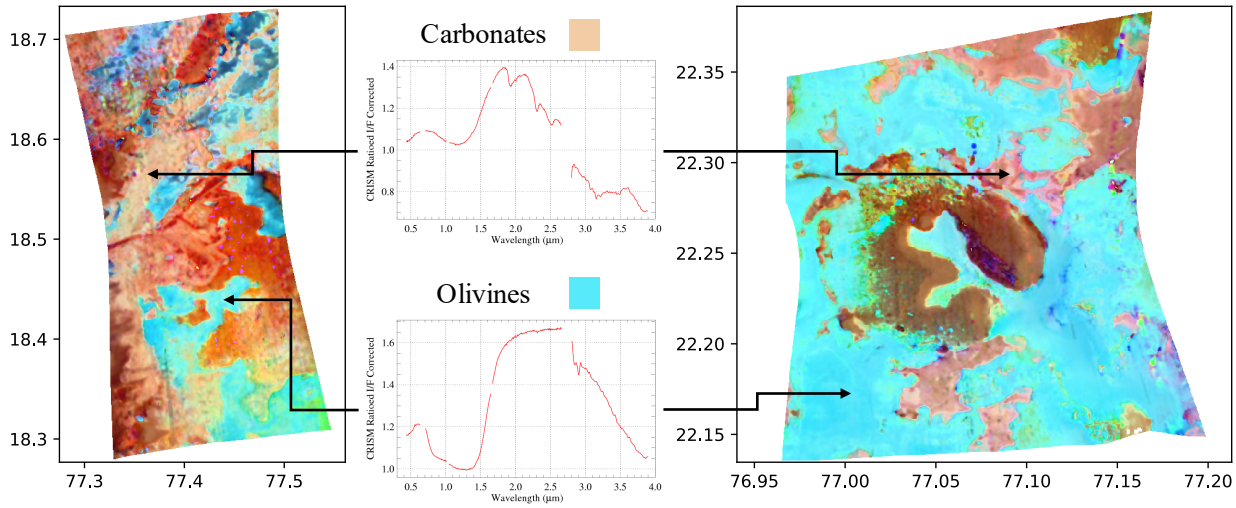

Figure 2: Spectroscopic maps of Martian regions at Jezero Crater (left) and Nili Fossae (right). Mineralogy is estimated from spectral signatures (center). Cream tones represent carbonates, whereas cyan indicates olivine presence.

2 RELATED WORK

Robots have become skilled at detecting objects under static conditions, goals and assumptions [11, 5, 12]. A few robots can go beyond detection to perform automatic science data analysis [13, 14]. However, they pursue static objectives that are fixed at the outset, such as detecting dust devils on Mars [15]. Some robots have used models based on Bayesian networks for mineral classification [12] or meteorite identification [11], but they also operate under predefined objectives that ignore the evolution of the robot's knowledge throughout the mission.

There are more adaptive models for autonomous robot exploration that rely on information-theoretic principles. Examples include Gaussian processes for binary terrain classification [5], ocean temperature mapping [8] and plant phenotyping [16]. Nevertheless, they are limited to scalar field mapping and are unable to reconstruct high resolution data. Recent efforts have developed adaptive algorithms with high-level science objectives [6, 17]. However, they work with highly discretized models and fail to exploit valuable information that is available from remote sensing data.

We are especially interested in previous work that combines low (remote) and high (*in situ*) resolution data for robotic exploration. Thompson *et al.* present a simple linear model that connects both, but actually downsamples the high resolution measurements and thus loses valuable information [18]. Candela *et al.* do the opposite by improving the details in low resolution data, but ignore important spatial correlations [19]. Foil and Wettergreen [20] and Thompson *et al.* [21] effectively utilize contextual information for surface classifications with spectroscopic data. However, informative robotic exploration needs a function to quantify and reward productivity, but in both cases it is either approximated [20] or ignored [21] due to computational difficulties.

Finally, recent work by Candela *et al.* manages to combine remote and *in situ* data for mapping spectroscopic data, derives an information-theoretic reward function for selecting the most valuable samples, and demonstrates feasibility with an autonomous rover carrying a spectrometer [10].

3 ACTIVE MAPPING

Planetary sciences rely on spectroscopic data for composition analysis. Each material reflects, emits, or absorbs light in a unique way throughout the wavelengths of the electromagnetic spectrum. The measured signals are called *spectra* and contain recognizable features or patterns that are used for rock and mineral identification [22]. *Spectroscopic maps* are spatial structures where each coordinate is associated with a full spectrum measurement and used to estimate composition. Figure 2 shows examples of spectroscopic maps at Jezero Crater and Nili Fossae.

We focus on the problem where a robotic explorer aims to learn the spatial distribution of spectra in the scene $\mathcal{M}$ in order to then map k mineral classes $c_1, c_2, \cdots, c_k$. We refer to this problem as *spatio-spectral regression*. We assume it can be done by combining two different types of measurements: low resolution *remote* spectra $x \in X \subset \mathbb{R}^m$, and high resolution *in situ* spectra $y \in$

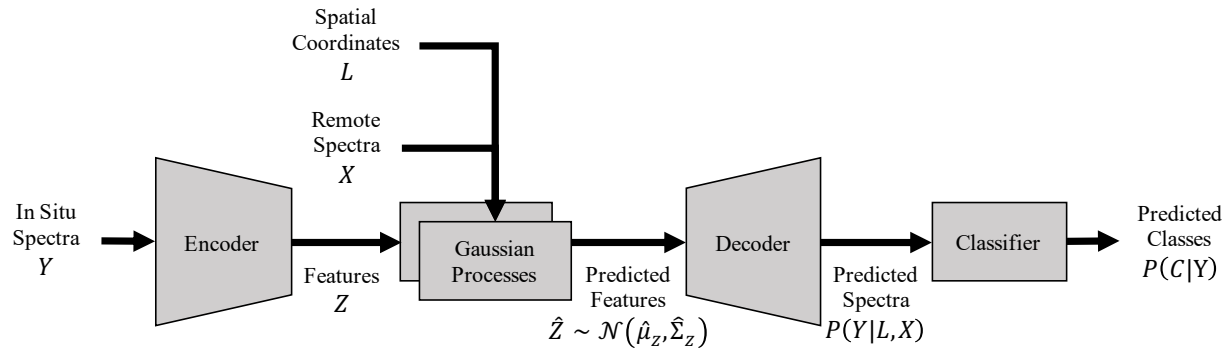

Figure 3: The active mapping model. It integrates feature extraction, regression, and classification. The rover collects in situ spectra $Y \subset \mathbb{R}^n$ and the encoder extracts their features $Z \subset \mathbb{R}^d$. Multiple Gaussian processes learn how to predict features from available remote spectra $X \subset \mathbb{R}^d$ and spatial coordinates $L \subset \mathbb{R}^2$. The decoder reconstructs high-resolution spectra from predicted features. The classifier estimates class membership C using predicted spectra.

$Y \subset \mathbb{R}^n$. Remote spectra are available beforehand for many spatial locations $l \in L \subset \mathbb{R}^2$ (e.g. latitude and longitude), whereas just a few *in situ* spectra can be collected by the rover.

We solve this problem by combining multiple machine learning algorithms. *Gaussian process regression* has been widely used in spatial statistics [23] and informative robotic exploration [5, 8, 16]. However, previous work focuses on mapping scalar fields (e.g. temperature), whereas our goal is to learn high-resolution spectroscopic maps. We tackle this issue with *feature extraction* techniques, which reduce the data dimensionality by deriving a subset of non-redundant features. Once a spectroscopic map is predicted, we utilize *classification* to estimate mineral composition. Figure 3 illustrates how these algorithms are combined in order to achieve active mineral mapping. The rest of this section describes feature extraction, Gaussian process regression, and classification in more detail.

3.1 Feature Extraction

In situ spectroscopic measurements have high resolution. Many channels and wavelengths are highly correlated (especially the ones that are adjacent to each other), allowing for the application of *dimensionality reduction* techniques. We use a variational autoencoder (VAE) [24], a neural network that performs nonlinear dimensionality reduction. It converts a set of high-resolution observations $y \in Y \subset \mathbb{R}^n$ into a set of lower dimensional features $z \in Z \subset \mathbb{R}^d$, where $d < n$. We specifically use a VAE because it learns a representation that resembles a standard multivariate normal distribution, i.e. $Z \sim \mathcal{N}_d(0, I_d)$, effectively normalizing and uncorrelating the features. The VAE is composed of two networks: an *encoder* that extracts the features, and a *decoder* that reconstructs high-resolution observations using the learned features. The architecture of the used VAE is similar to the one presented by Candela *et al.* [19]. The VAE ignores spatial information.

3.2 Gaussian Processes for Spatio-Spectral Regression

The robot uses Gaussian processes (GPs) [25] for spatio-spectral regression, that is, to learn the spatial distribution of spectra throughout the scene. GPs are typically used for mapping scalar values, but our problem involves multivariate regression. We simplify by using GP regression to learn the distribution of low dimensional features Z instead. Moreover, dimensionality reduction with a VAE uncorrelates the learned feature representation, allowing for the utilization of d independent GPs.

We next provide a brief explanation in regard to our specific GP regression model. If needed, extensive and canonical documentation can be found in [25]. Formally, we define an input vector that concatenates spatial coordinates and remote measurements as $v = [l, x] \in V \subset \mathbb{R}^{2+m}$, similarly as Thompson [5]. We assume there exists a latent function $f^i : \mathbb{R}^{2+m} \to \mathbb{R}$ that maps an input v to each feature $z^i = f^i(v) + \epsilon_i$, where $i = 1, 2, \ldots, d$. Each GP learns a distribution over the values that f^i can take. A GP is defined by a mean function μ^i and a covariance function K^i_θ, i.e. $f^i(v) \sim \mathcal{GP}(\mu^i(v), K^i_\theta(v, v'))$. We assume that the mean is zero because of the way features are normalized by the VAE. For the covariance matrix, we rely on the widely used *squared exponential kernel*. We define an anisotropic kernel that distinguishes between spatial and spectral dimensions:

$$K^i_\theta(v, v') = \theta^i_0 \exp\left(-\frac{\|l - l'\|_2^2}{2(\theta^i_l)^2} - \frac{\|x - x'\|_2^2}{2(\theta^i_x)^2}\right), \qquad (1)$$

where $\theta^i = \left[\theta_0^i, \theta_l^i, \theta_x^i\right]$ are the kernel hyperparameters for each GP. Additionally, we utilize the GP variant for noisy observations, and thus use an additional hyperparameter σ^i_{noise} . The GP hyperparameters are estimated by maximizing the log-likelihood of the observed data as shown in [25].

3.3 Mineral Classification from Spectra

Mineral composition is automatically estimated via classification. A predicted high-resolution spectrum y is compared against a set of well-known spectra, one for each mineral class $c_1, c_2, \cdots, c_k$. This set is known as a *spectral library*. In this paper we use a Gaussian classifier that estimates class membership using Bayes's theorem as follows:

$$p(c|y) = \frac{p(c)p(y|c)}{\sum_{c'} p(c')p(y|c')}, \tag{2}$$

where $p(c|y)$ is the probability of class membership, $p(c)$ is the prior distribution, and the likelihood follows a Gaussian distribution, i.e. $p(y|c) \sim \mathcal{N}(\mu, \Sigma)$. Each class c_j is defined by a mean μ_j and a covariance Σ_j.

4 INFORMATIVE EXPLORATION

The rover aims to collect meaningful science measurements: the ones that better explain and reconstruct the scene. Both in information theory and Bayesian experimental design, information-driven action selection can be formulated as the minimization of *posterior entropy*, which measures the uncertainty of a variable of interest after collecting new information [4, 6, 5, 16, 7]. Here the variable of interest is the spatial distribution of features throughout a spectroscopic map, which we assume is composed of a large, yet finite set of points $\mathcal{M} = \{v_1, \ldots, v_p\}$. The new information is given by the *in situ* coordinates and measurements collected by the robot, i.e. $\mathcal{P} = \{[l_1, y_1], \ldots, [l_k, y_k]\}$. The uncertainty of the model is quantified using Shannon entropy, which is additive for independent sources. Since the model consists of d independent GPs, the entropy of the map is additive for features and given by the following expression [4, 7]:

$$H(\mathcal{M}|\mathcal{P}) = \frac{1}{2}\sum_{i=1}^{d} \log\left((2\pi e)^{|\mathcal{M}|}\,|\hat{\Sigma}^i_{\mathcal{M}|\mathcal{P}}|\right), \tag{3}$$

where $\hat{\Sigma}^i_{\mathcal{M}|\mathcal{P}}$ is the predicted covariance of each feature z^i throughout the entire map $\mathcal{M}$ after being updated with the *in situ* samples $\mathcal{P}$ [25]. Then, an informative planner will solve an optimization problem where the path should minimize the posterior entropy of the map:

$$\min_{\mathcal{P}} H(\mathcal{M}|\mathcal{P}) \quad \text{subject to} \quad Cost(\mathcal{P}) \leq Budget. \tag{4}$$

There are many path planners that could be used for this problem. Some methods assume independence between sampling locations, but this is usually an oversimplification in informative exploration scenarios [26]. There are near-optimal greedy algorithms that work well on Gaussian processes [7, 16]. Other approaches may be computationally intensive, but potentially closer to optimality, such as branch and bound techniques planning techniques [8]. Finally, we are especially interested in Monte Carlo tree search (MCTS) planners that have been applied to geologic exploration scenarios [9, 17].

5 MARS SIMULATION STUDY

We demonstrate the performance of our approach in a simulation study. It involves spectroscopic investigations and mineral mapping at two locations of interest on Mars: Jezero Crater and Nili Fossae.

This study uses or simulates data from the following three spectrometers: the Mars Express High-Resolution Stereo Camera (HRSC) [27], the Mars 2020 Perseverance rover multispectral imager MastCam-Z [28], and the MRO Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) [29].

HRSC serves as the source of remote data because it has global coverage of Mars, but its resolution is insufficient for composition analysis. It has an approximate resolution of 50 m/pixel and 4 color channels: blue 440 nm, green 530 nm, red 750 nm, and infrared 970 nm. Image h5270_0000 is used for both sites.

MastCam-Z serves as the source of *in situ* data, allowing for mineral identification in the visible and near-infrared wavelengths (400 - 1100 nm). MastCam-Z consists of 2 cameras (left and right) with a total of 20 bandpass filters (channels). Since MastCam-Z data is yet to be acquired, we simulate rover measurements using MastCam-Z spectral response information in conjunction with data from CRISM.

CRISM is an imaging spectrometer that provides high-resolution spectra (6.5 nm/channel) in a wide spectral range (436 - 3897 nm). It has a spatial resolution of 18-36 m/pixel. The following CRISM images are used in this study (Figure 2): HRL000040FF (Jezero Crater), and FRT00003E12 (Nili Fossae).

Data from these spectrometers is prepared as follows. The HRSC and CRISM products are spatially aligned with manually-selected ground control points. They are co-registered using a first degree polynomial warping transformation, and then resampled to 50 m/pixel. CRISM data is preprocessed and ratioed using the standard procedure in [2]. Afterwards, the empirical line method [30] is used to find the correspondence

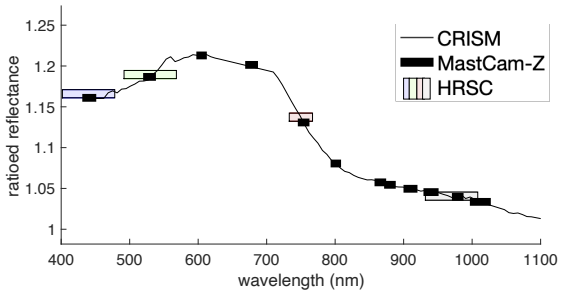

Figure 4: A mineral (Fe olivine) as measured by the Mars instruments CRISM, MastCam-Z, and HRSC. Their channels and respective bandwidths are shown.

between the HRSC and CRISM reflectance values. MastCam-Z is estimated by convolving its spectral response profiles [28] with CRISM data. A subset of 12 channels is used since 8 of the filters either overlap (i.e. too redundant) or have very wide bandwidths. Figure 4 shows an example of this curated data set.

This study relies on the CRISM spectral library for mineral classification since it contains many of the minerals that have been identified on Mars [2]. In this paper we focus on carbonates since they have been found on both sites and are of scientific interest due to their fossil preservation properties [2, 31]. We also consider olivines and silicas. Specifically, our classifier uses 5 minerals from the CRISM spectral library: Fe/Ca-carbonate, Mg-carbonate, Fe-olivine, Mg-olivine, and hydrated silica.

6 EXPERIMENTS AND RESULTS

The experiments are designed to evaluate the performance of both the learning and exploration strategies in terms of science productivity.

6.1 Experimental Setup

Three hundred rover traverses are simulated at each of the two sites: Jezero Crater and Nili Fossae. For each, we utilize a 3×3 km subregion that is mineralogically diverse. We represent the exploration space as an 8-connected grid with a step size of 150 meters (3 pixels). Starting locations are evenly spaced throughout each subregion; end goals are not specified. We impose a constraint of 20 samples per traverse.

We compare the performance of three informative path planning algorithms:

- **Random**: sequentially samples a random neighboring location until the sampling budget is exhausted. This is a *science-blind* baseline.
- **Greedy**: sequentially samples the best neighboring location using a one-step lookahead. This is a *myopic* exploration strategy.
- **MCTS**: the Monte Carlo tree-search planner by Kodgule *et al.* [9] using a four-step lookahead. This is a *non-myopic* path planner.

We use three metrics to evaluate the performance of the planners. For normalization purposes, we compute the averages with respect to the total number of points in the map. The first metric is the posterior *entropy* (Equation 3), which is *directly* minimized by the planners (Equation 4) and is calculated without a ground truth. The second metric is the *reconstruction error* of spectra throughout the scene in terms of root mean squared error (RMSE). It should be *indirectly* minimized by the planners since it always requires a basis for comparison. The third metric is the *Kullback-Leibler divergence* (KLD); it is a smooth function that measures the difference between the real and the predicted class probabilities.

In this study, both the VAE and the GPs (Section 3) are trained with the data that is withheld from the experiments. L2 normalization is applied to both HRSC and MastCam-Z spectra to allow the model to focus on spectral features rather than albedo values. The VAE extracts features from MastCam-Z spectra and encodes them into a space with dimensionality $d = 3$; a value we find to work well. The model consists of 3 independent GPs that are pre-trained with the same data set, and later fine-tuned online in order to better adapt to incoming data. The Gaussian classifier consists of the following parameters that we find to work well in practice: a uniform prior distribution, means corresponding to the 5 selected minerals from the CRISM spectral library, and isotropic (spherical) covariances $\Sigma = (0.02)^2 I_n$. We assume that the ground truth classes are given by the output of the Gaussian classifier when applied to simulated MastCam-Z spectra, which is also assumed to be the ground truth *in situ* spectra.

Training and simulation are performed using a laptop computer with an Intel i7 processor (2.9 GHz quad-core) and 16GB of memory. Each waypoint is computed within just a few seconds or minutes, depending on the complexity of the path planner.

6.2 Results

We first discuss and show qualitative results. Figure 5 includes an example of a rover traverse at Nili Fossae together with its corresponding mineral mapping process. Mineral signatures are successfully identified and extrapolated throughout the entire scene. We observe that just a few samples (20) are sufficient to map the mineralogy of most of the 3×3 km area.

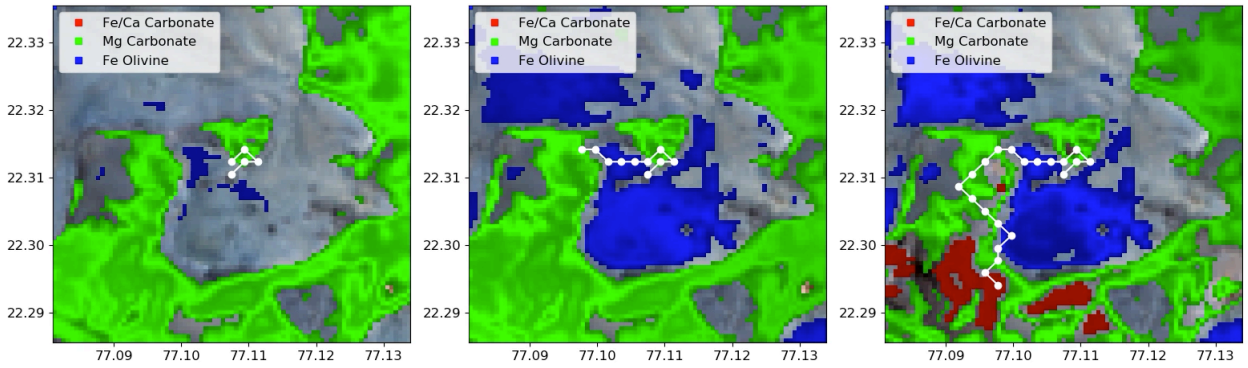

Figure 5: Example of the mineral mapping process at Nili Fossae. Left: the rover identifies and extrapolates two mineral signatures. Center: the rover collects more samples and updates the map. Right: the rover discovers a third mineral. This example only shows points with a mineral probability greater than 50%.

We now analyze metric correlation throughout the simulations. Figure 6 shows error vs. entropy scatter plots for Jezero Crater and Nili Fossae with correlation coefficients of 0.8855 and 0.8262, respectively. These values indicate a positive correlation between HRSC and CRISM data, and confirm that entropy is a suitable objective function for spectroscopic mapping. Figure 7 shows KLD vs. error plots with even higher correlations: 0.9806 and 0.9609. This makes sense since classification accuracy (KLD) directly depends on the quality of the predicted high-resolution spectra.

We then evaluate the performance of the three planners. The corresponding plots are shown in Figure 8, where all results have statistically significant differences ($\alpha = 0.05$). In all cases, entropy, reconstruction error, and KLD show decreasing trends as more samples are collected. Error and divergence converge faster than entropy. It is clear that Random is the worst planner. We confirm that MCTS, the algorithm with the farthest planning horizon, performs best at the end. Note that Greedy outperforms MCTS during the first few samples; this is to be expected because Greedy selects points that immediately provide high rewards, whereas MCTS computes a long-term strategy.

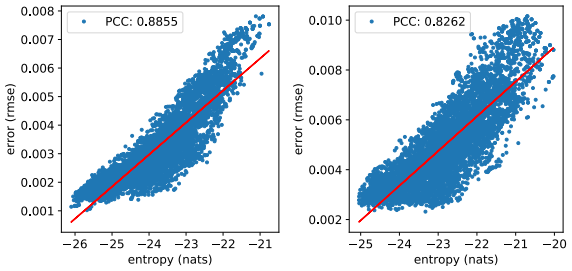

Figure 6: Average reconstruction error vs. entropy at Jezero Crater (left) and Nili Fossae (right).

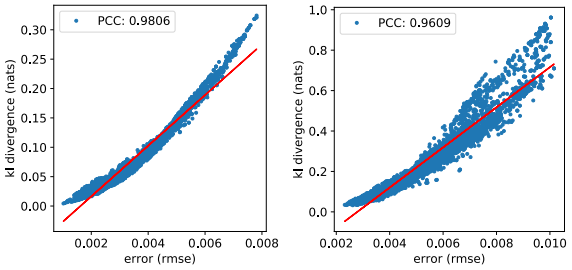

Figure 7: Average KLD vs. reconstruction error at Jezero Crater (left) and Nili Fossae (right).

7 CONCLUSIONS

This paper presents an approach to improve science productivity in autonomous rover exploration. We describe an active learning model for wide-area mineral mapping that combines remote and *in situ* measurements. It extrapolates mineral signatures from just a few rover samples in order to densely map an explored scene. This is done by fusing feature extraction, Gaussian process regression, and mineral classification. Furthermore, exploration and productivity are improved by incorporating notions from information theory and non-myopic path planning.

The presented Mars simulation study confirms the results that were also found in previous work at Cuprite, Nevada. Entropy and reconstruction error are strongly correlated, showing that entropy is a suitable objective function for scene mapping. The model's entropy, reconstruction error, and Kullback-Leibler divergence present decreasing trends as a function of sampling. We also observe that non-myopic planning consistently outperforms simpler exploration strategies such as myopic planning and random sampling.

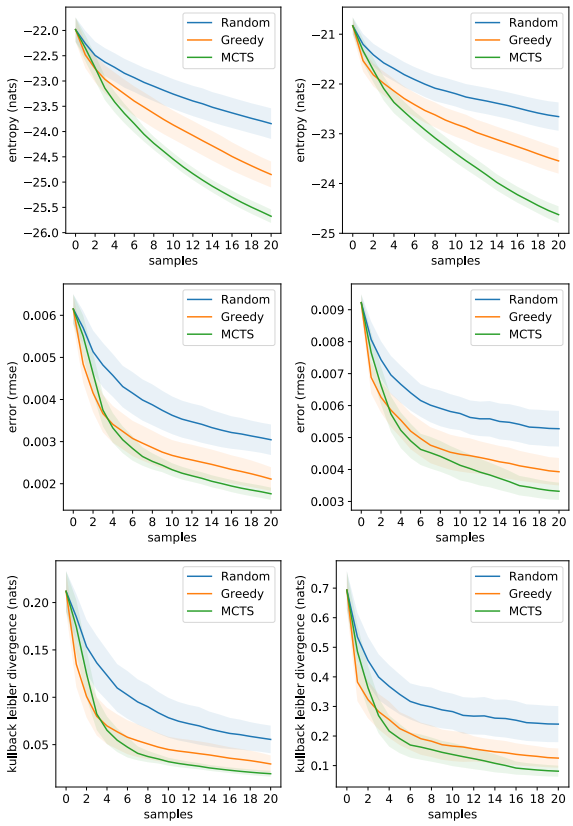

Figure 8: Performance plots and 1-sigma error bars for Jezero Crater (left column) and Nili Fossae (right column). Average entropy (top row), reconstruction error (middle row), and KLD (bottom row) as a function of collected samples per traverse. All results have a statistically significant difference ($\alpha = 0.05$).

Future investigations will address rover traversability, specifically the trade-off between safety and science productivity in heterogeneous terrains, as well as the potential of mineralogy mapping for terrain understanding. Finally, we will apply these methods to other scientific endeavors that could also benefit from autonomous robotic exploration, such as maritime or agricultural mapping.

Acknowledgment

This research was supported by the National Science Foundation National Robotics Initiative Grant #IIS-1526667, and by the NASA Solar System Exploration Virtual Institute and the Planetary Sciences Institute project TREX: Toolbox for Research and Exploration (award 80ARC017M0005).

References

[1] Gaines D, Russino J, Doran G, Mackey R, Paton M, Rothrock B, Schaffer S, Agha-Mohammadi AA, Joswig C, Justice H, Kolcio K, Sawoniewicz J, Wong V, Yu K, Rabideau G, Anderson R and Vasavada A (2018) Self-Reliant Rover Design for Increasing Mission Productivity. In: *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*.

[2] Viviano-Beck CE, Seelos FP, Murchie SL, Kahn EG, Seelos KD, Taylor HW, Taylor K, Ehlmann BL, Wiseman SM, Mustard JF and Morgan MF (2014) Revised CRISM spectral parameters and summary products based on the currently detected mineral diversity on Mars. In: *Journal of Geophysical Research: Planets*, 119(6):pp.1403–1431.

[3] Bishop JL (2018) *Chapter 3 - Remote Detection of Phyllosilicates on Mars and Implications for Climate and Habitability*. In: NA Cabrol and EA Grin (Editors), *From Habitability to Life on Mars*, pp.37 – 75. Elsevier.

[4] Cover TM and Thomas JA (2006) *Elements of Information Theory*. Wiley- Interscience, Hoboken, NJ, 2nd edition.

[5] Thompson DR (2008) *Intelligent Mapping for Autonomous Robotic Survey*. Doctoral Thesis, Carnegie Mellon University, Pittsburgh, PA.

[6] Candela A, Thompson D, Dobrea EN and Wettergreen D (2017) Planetary robotic exploration driven by science hypotheses for geologic mapping. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp.3811–3818.

[7] Krause A, Singh AP and Guestrin C (2008) Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. In: *Journal of Machine Learning Research*, 9:pp.235–284.

[8] Binney J and Sukhatme GS (2012) Branch and bound for informative path planning. In: *2012 IEEE International Conference on Robotics and Automation*, pp.2147–2154.

[9] Kodgule S, Candela A and Wettergreen D (2019) Non-myopic Planetary Exploration Combining In Situ and Remote Measurements. In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp.536–543.

[10] Candela A, Kodgule S, Edelson K, Vijayarangan S, Thompson DR, Noe Dobrea E and Wettergreen D (2020) Planetary Rover Exploration Combining Remote and In Situ Measurements for Active Spectroscopic Mapping. In: *IEEE International Conference on Robotics and Automation (ICRA)*, pp.5986–5993.

[11] Pedersen L, Wagner M, Apostolopoulos D and Whittaker WR (2001) Autonomous robotic meteorite identification in Antarctica. In: *IEEE International Conference on Robotics and Automation*, pp.4158–4165 vol.4.

[12] Gallant M, Ellery A and Marshall J (2011) Science-influenced mobile robot guidance using Bayesian Networks. In: *Canadian Conference on Electrical and Computer Engineering*, IEEE, pp.1135–1139.

[13] Estlin TA, Bornstein BJ, Gaines DM, Anderson RC, Thompson DR, Burl M, Castaño R and Judd M (2012) AEGIS Automated Science Targeting for the MER Opportunity Rover. In: *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(3).

[14] Chien S, Sherwood R, Tran D, Cichy B, Rabideau G, Castano R, Davis A and Boyer D (2005) Using autonomy flight software to improve science return on earth observing one. In: *Journal of Aerospace Computing, Information, and Communication*, 2:pp.196–216.

[15] Castano A, Fukunaga A, Biesiadecki J, Neakrase L, Whelley P, Greeley R, Lemmon M, Castano R and Chien S (2006) Autonomous Detection of Dust Devils and Clouds on Mars. In: *2006 International Conference on Image Processing*, pp.2765–2768.

[16] Kumar S, Luo W, Kantor G and Sycara K (2019) Active Learning with Gaussian Processes for High Throughput Phenotyping. In: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, International Foundation for Autonomous Agents and Multiagent Systems, Montreal QC, Canada, pp.2078—-2080.

[17] Arora A, Fitch R and Sukkarieh S (2017) An approach to autonomous science by modeling geological knowledge in a Bayesian framework. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp.3803–3810.

[18] Thompson DR, Wettergreen D, Foil G, Furlong M and Kiran AR (2015) Spatio-spectral exploration combining in situ and remote measurements. In: *Twenty-Ninth AAAI Conference on Artificial Intelligence*, pp.3679–3685.

[19] Candela A, Thompson DR and Wettergreen D (2018) Automatic Experimental Design Using Deep Generative Models Of Orbital Data. In: *International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*.

[20] Foil G and Wettergreen D (2018) Efficiently Sampling from Underlying Models. In: *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania.

[21] Thompson DR, Candela A, Wettergreen DS, Noe Dobrea E, Swayze GA, Clark RN and Greenberger R (2018) Spatial Spectroscopic Models for Remote Exploration. In: *Astrobiology*, 18(7):pp.934–954.

[22] Clark RN (1999) Spectroscopy of rocks and minerals, and principles of spectroscopy. In: *Manual of Remote Sensing*, 3:pp.3–58.

[23] Cressie NA (1993) *Statistics for Spatial Data*. John Wiley & Sons, Inc.

[24] Kingma DP and Welling M (2013) Auto-Encoding Variational Bayes. In: *arXiv preprint arXiv:1312.6114*.

[25] Rasmussen CE and Williams CKI (2006) *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge, Massachusetts.

[26] Yilmaz NK, Evangelinos C, Lermusiaux PFJ and Patrikalakis NM (2008) Path Planning of Autonomous Underwater Vehicles for Adaptive Sampling Using Mixed Integer Linear Programming. In: *IEEE Journal of Oceanic Engineering*, 33(4):pp.522–537.

[27] Jaumann R et al. (2007) The high-resolution stereo camera (HRSC) experiment on Mars Express: Instrument aspects and experiment conduct from interplanetary cruise through the nominal mission. In: *Planetary and Space Science*, 55(7-8):pp.928–952.

[28] Bell JF, Maki JN, Mehall GL, Ravine MA, Caplinger MA and Mastcam-Z Team (2016) Mastcam-Z: Designing a Geologic, Stereoscopic, and Multispectral Pair of Zoom Cameras for the NASA Mars 2020 Rover. In: *3rd International Workshop on Instrumentation for Planetary Mission*, volume 1980, p.4126.

[29] Murchie S et al. (2007) Compact Connaissance Imaging Spectrometer for Mars (CRISM) on Mars Reconnaissance Orbiter (MRO). In: *Journal of Geophysical Research E: Planets*, 112(5):pp.1–57.

[30] Smith GM and Milton EJ (1999) The use of the empirical line method to calibrate remotely sensed data to reflectance. In: *International Journal of Remote Sensing*, 20(13):pp.2653–2662.

[31] Horgan BH, Anderson RB, Dromart G, Amador ES and Rice MS (2020) The mineral diversity of Jezero crater: Evidence for possible lacustrine carbonates on Mars. In: *Icarus*, 339:p.113526.