INTELLIGENT MAPS FOR AUTONOMOUS KILOMETER-SCALE SCIENCE SURVEY

David R. Thompson and David Wettergreen

The Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh PA 15213, USA
drt@ri.cmu.edu, dsw@ri.cmu.edu

ABSTRACT

We present a new approach for site survey by autonomous surface robots. In our method the agent constructs an *intelligent map*, a multi-scale model of the explored environment incorporating *in situ* and remote sensing data. The agent learns the model's parameters on the fly and exploits its predictions to guide adaptive navigation and sampling. In this manner the agent can respond appropriately to novel correlations, resource constraints and execution errors. Rover tests at Amboy Crater, California demonstrate improved performance over non-adaptive strategies for a geologic survey task.

1. INTRODUCTION

Teleoperation of planetary exploration robots is often infeasible due to the infrequency and low bandwidth of communications with the remote agent. Under these circumstances onboard analysis can improve the quality of returned science data. Autonomous agents can explore continuously, identify informative features for measurement, and select the best data products to transmit [1]. Onboard science autonomy has enabled new operational modes like selective data return by the EO-1 spacecraft [2] and automatic dust devil detection by the Mars Exploration Rovers [3]. This previous work generally involves data analysis that detects one or more specific predefined features of interest.

Here we consider a new application for onboard data understanding: a site survey task in which an agent characterizes a large area using a small number of samples. The agent must allocate observations intelligently to reduce redundancy and make best use of limited time and energy resources. Because agent's goal is to characterize the whole site, the value of each new observation hinges on previous knowledge gleaned from other measurements. In other words, the site survey application requires a context-sensitive formulation of data's science value. Agents must respect correlations in data between samples at different surface locations or between surface and remote sensing data.

Figure 1. The rover Zoë at Amboy Crater approaching a boundary between basalt and sediment surface units.

Here we treat site survey as an experimental design task, selecting observations that provide the most information about a spatial model of the environment. A Gaussian process model represents correlations between remote sensing and *in situ* measurements. This *intelligent map* improves exploration in several respects. The map itself is a useful data product; it is a compressed description of the traverse. Moreover, model predictions facilitate context-sensitive sampling and navigation decisions. Prediction uncertainty determines the information value of candidate observations. The agent learns model parameters during exploration so that these science value assessments adapt to reflect unanticipated trends and correlations.

This work applies the intelligent map approach to a surficial site survey. Here a rover travels within a predefined "exploration corridor" toward an end-of-day goal location while constructing maps of surface features (Figure 2). We describe mechanisms for model learning and inference, and show how prediction uncertainty can inform adaptive sampling. Finally we describe tests at Amboy Crater, California that demonstrate kilometer-scale exploration. Here a rover builds surficial maps by learning and exploiting correlations between *in situ* reflectance spectroscopy and remote sensing data (Figure 3).

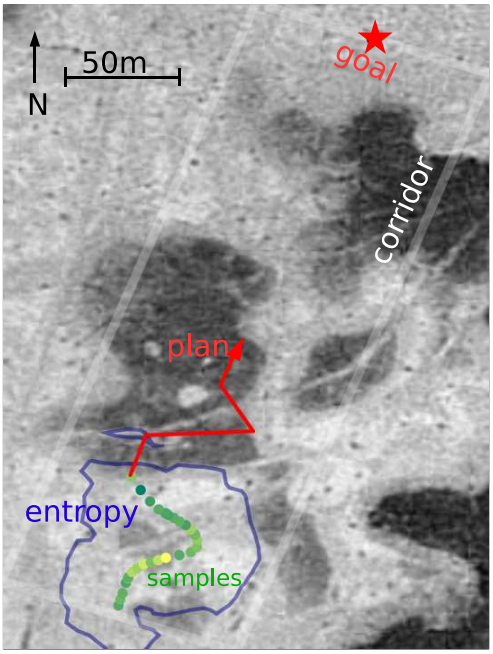

Figure 2. Overflight imagery of the field site. Annotations show the rover status during an autonomous traverse. Colored dots show spectral samples collected along the rover path. The blue line indicates one isocontour of the marginal prediction entropy, whose significance is explained in Section 3. The red line shows waypoints in the current navigation plan.

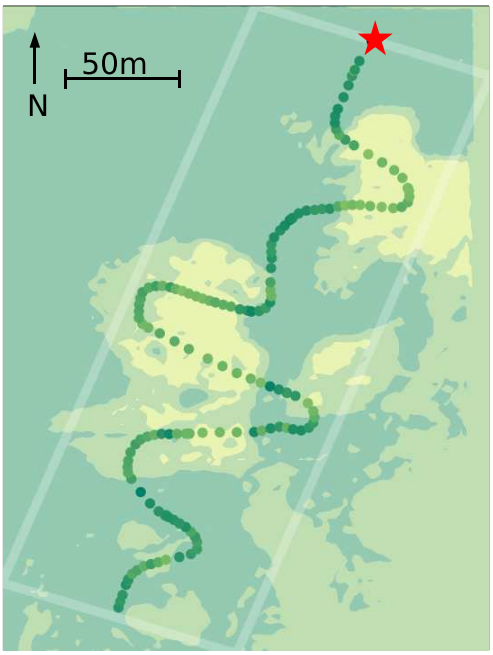

Figure 3. Geologic map of surface material generated by the rover from the traverse of Figure 2. Gaussian process inference extrapolates from sparse samples by leveraging remote sensing and proximity correlations. The rover targets areas for which predictions are most uncertain.

2. GAUSSIAN PROCESS MODELS

Briefly, a Gaussian process [4] is a nonparametric Bayesian scheme for modelling a regression of independent variables $X = \{x_1, x_2, \ldots x_n\}$, $x_i \in \mathbb{R}^d$ onto scalar dependent observations $Y = \{y_1, y_2, \ldots, y_n\}$, $y_i \in \mathbb{R}$. In our site survey domain x_i is a vector of sample site attributes including spatial location and any associated remote sensing data. For each site x_i we transform the latitude and longitude position (l_{i0}, l_{i1}) along with m remote image channels $(c_{i0}, \ldots, c_{im})$ to a common scale, and append them to produce the attribute vector:

$$x_i = (l_{i0}, l_{i1}, r_{i0}, \ldots, r_{im}) \quad (1)$$

This is the input to an unknown function f mapping site attributes onto an observation y_i with additive Gaussian noise: $y_i = f(x_i) + \epsilon$. For our site survey task the observation y is a single scalar value describing the surface material observed at a sample site. The Gaussian process model gives a distribution over underlying functions f such that its evaluation at any set of potential sample sites has a multivariate Gaussian distribution. Evaluating marginal *maximum a posteriori* estimates $\hat{f}(x)$ at many unobserved locations yields a surficial map of the desired extent and resolution (Figure 3).

2.1. Stationary Covariance Functions

The covariance function $\kappa(x_i, x_j)$ determines a prior over f in terms of correlations between input points x_i and x_j. Suppose the explorer agent collects observations at sample sites $X = \{x_1, \ldots, x_n\}$ with noise variance s^2. This yields a positive definite covariance matrix C with additive noise on the diagonal:

$$C_{i,j} = \kappa(x_i, x_j) + s^2\delta(x_i, x_j) \quad (2)$$

These measurements induce a posterior distribution over a future observation at x_{n+1}. Define $q = \kappa(X, x_{n+1})$ to be a vector consisting of the covariance function evaluated between each training point and the new sample site. The posterior mean $\hat{\mu}$ and variance $\hat{\sigma}^2$ of the new observation is given by the following expression [4] involving the vector of training outputs Y:

$$\hat{f}(x_{n+1}) = \hat{\mu} = q^T C^{-1} Y \quad \text{and} \quad \hat{\sigma}^2 = q^T C^{-1} k \quad (3)$$

In this work we use a popular *squared exponential* covariance function with hyperparameters θ_1, θ_2, and a bandwidth w_k for each input dimension $k = 1, \ldots, d$:

$$\kappa(x_i, x_j) = \theta_1 + \exp\{-\theta_2 \sum_{k=1}^{d} w_k (x_{ik} - x_{jk})^2\} \quad (4)$$

Bandwidth hyperparameters reflect smoothing along each dimension of the input data. In other words, they describe the correlation between two measurements based on their location and remote sensing attributes. If the

geologic class of surface material is strongly correlated with a particular remote sensing channel, a good exploration policy would visit sites that are diverse with respect to that remote sensing band. Alternatively, if spatial location is a good predictor of geologic content then the appropriate exploration policy evenly covers the physical space. Different bandwidth parameters might place the optimal policy between these extremes.

Appropriate hyperparameters are difficult to specify in advance so it is desirable to learn them on the fly during exploration. Convention favors two alternative estimation approaches: Markov Chain Monte Carlo, or gradient descent of the hyperparameter likelihood. The latter is susceptible to local minima, but these were rare in our application so we opted for the speed of conjugate gradient descent. The data likelihood $L(w,k)$ for our covariance function has the following general form [4]:

$$L(w,k) = \quad -\tfrac{1}{2}\log\det C - \tfrac{1}{2}Y^T C^{-1} Y - \tfrac{n}{2}\log 2\pi \quad (5)$$

We reparameterize the likelihood so that the it respects known isotropies in the input dimensions. For example, bandwidth parameters associated with the two spatial dimensions should be equivalent. After finding the hyperparameter gradient we use the Fletcher-Reeves formula to identify conjugate gradient directions. We perform a line search to maximize $L(w,k)$ and recompute the conjugate gradient from the new hyperparameter values. This method converges in a several iterations. We avoid local maxima by restarting the optimization at random locations in the hyperparameter space.

2.2. Nonstationary Covariance Functions

The covariance function of equation 4 is *stationary*; two observations' covariance depends only on their relative locations in the input space. This is a poor model for surficial geologic units which often include homogeneous areas and discrete boundaries. Figure 1 shows typical terrain from the field expedition which exhibits a sharp discontinuity between regions of sediment and basaltic lava. These contacts are better described by *nonstationary* covariance where the rate of smoothing varies over the input space [5].

Our experiments model nonstationary covariance with the method of Pfingsten *et. al* [6]. The approach augments the input space with the prediction of a stationary regression model. The new input space separates training data with different output values along the additional dimension so that their covariance is nonstationary with respect to the original input. In our implementation the latent regression model is also a Gaussian process. We train the latent model on collected data to produce an initial prediction estimate $\hat{f}(x)$ for each datapoint. We then scale this prediction to the $[0,1]$ interval and append it back to the input space. This procedure transforms the original independent variable x_i to yield a new x_i'. The attribute vector of a sample site becomes:

$$x_i' \quad = \quad (l_i^0, l_i^1, r_i^0, \ldots, r_i^m, 1 + exp\{\hat{f}(x_i)\}^{-1}) \quad (6)$$

We then perform a second round of inference using a Gaussian Process model trained on the transformed input vectors x_i'. This augmented model has an additional free bandwidth parameter to reflect its extra input dimension.

Figure 4 shows the impact of this transform on synthetic data from a noisy step function. A stationary covariance function forces hyperparameters to compromise between estimating smooth and discontinuous regions. The resulting model undersmoothes at the periphery and oversmoothes at the discontinuity (Figure 4 left). The Gaussian process with augmented inputs produces a superior estimate (Figure 4 right). Note that the stationary covariance function's prediction variance is high near the periphery where the data is sparse. This contrasts with the nonstationary case where prediction uncertainty is high near the step discontinuity. The nonstationary prediction variance reflects uncertainty about observations near the area of high change. Discrete transitions lead naturally to a "boundary following" policy in which the explorer attempts to localize unit borders. Alternatively, more graded unit transitions favor a strategy with greater spatial coverage.

3. BAYESIAN EXPERIMENTAL DESIGN

The explorer agent can leverage Gaussian process predictions to improve its exploration plan. We employ a *Bayesian experimental design* strategy in which the agent selects measurements that minimize posterior uncertainty about the explored environment. For our specific site survey application the values of interest consist of the regression function $f(x)$ at a hypothetical set of unmeasured locations. The agent chooses observations in order to reduce the uncertainty over these sample locations as quantified by their Shannon entropy.

Shewry and Wynn provide assumptions under which this spatial design problem reduces to one of maximizing the entropy of the agent's observations [7]. They require that the underlying process be finite and observations have uniform noise which is independent of the sampling strategy. For any set of potential measurement points Q the entropy $H(Q)$ is a fixed constant. If we aim to choose an observation subset A to reduce posterior entropy over unobserved sites $Q \setminus A$, the objective decomposes:

$$E[H(Q \setminus A \mid A)] \quad = \quad H(Q) - H(A) \quad (7)$$

This leads to the strategy of maximum entropy sampling [7, 8, 9]. Since $H(Q)$ is a fixed property of the environment, optimal data collection maximizes the entropy $H(A)$ of the observations. Maximum entropy sampling reflects an intuitive idea that one should collect data where the result is most uncertain. The approach has a long history of use in geostatistics and Gaussian process models, but it is worth noting that studies have presented alternative measures of information gain such as the mutual information criterion of Guestrin *et al.* [10]. In our task we found that maximum entropy sampling provided reasonable performance and computational efficiency.

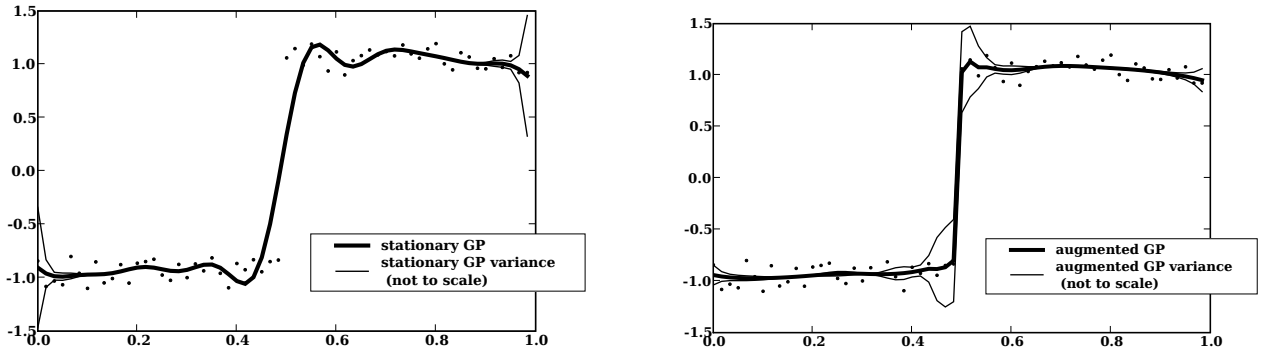

Figure 4. Stationary Gaussian Process regression oversmoothes at a step discontinuity (left). Prediction variance depends only on proximity to training data. Augmented inputs produce nonstationary data-dependent covariance structure (right).

Exploration rovers performing site survey tasks must characterize large explored environments using only a few samples. They can benefit from adapting their experimental design during execution. This compensates for inaccurate initial estimates of covariance function hyperparameters. Moreover, adaptive designs can exploit resource surpluses or recover from shortages whenever execution errors modify the original plan. At regular intervals our agent reestimates the covariance function using all collected data and creates a new maximum entropy observation plan for the remaining resource budget. This is *Bayesian adaptive exploration* [11], which has resurfaced in such varied fields as robotic exploration and mapping [12, 13, 14], active learning [15], and more broadly in classical experimental design [16, 17].

3.1. The Corridor Exploration Task

We formulate adaptive resource allocation as an optimization over sequences of navigation and data collection actions. The most general solution would consider the potential of future observations to change the Gaussian process model itself. This is tantamount to a Partially-Observable Markov Decision Process (POMDP). The general POMDP exploration task is intractable, but recent work with off-line optimization has made progress using approximate solutions and cleverly-factorized state spaces [18, 19]. However, even if we assume the model is correct, identifying a maximum-entropy subset of sampling locations is NP-hard [9]. Branch-and-bound approximations exist for sample locations on a connected graph [9, 20], but real-time optimization of arbitrary data collection and navigation action sequences remains impractical for continuous rover exploration domains.

Here we consider a constrained "corridor exploration" task in which the agent moves forward inside a predefined exploration area toward a specific end-of-day location. The agent builds a surficial map based on samples collected at regular intervals along the traverse path. Figure 5 illustrates the scenario: a rover travels the path which maximizes observation entropy while respecting the available time budget.

This corridor exploration task simplifies the general problem in several useful ways. First, the agent incurs no explicit time or resource cost for data collection. This is consistent with observations like images or reflectance spectra that are relatively easy to collect. The simplification also applies whenever data collection occurs automatically at a fixed rate. Incorporating data collection cost into movement cost reduces the joint optimization problem to one of path planning.

Corridor exploration also restricts the form of the utility function. General utility functions may have multiple terms balancing expected information gain against the cost of time and other resources [14]. Here our only significant resource is time — for a set of drive actions D the time cost $C(D)$ is proportional to the total path length. With some abuse of notation we write the joint entropy of observations along this path as $H(D)$. In general the designer could assign weighting coefficients to both terms and adjusts them through trial and error in order to achieve desired behavior.

In the corridor exploration task movement cost is implicit in the requirement that the robot be able to reach the end-of-day goal location. This is a discontinuous utility function: cost is zero until the budget is exhausted, at which point it becomes infinite. The path planning algorithm need only explore the space of paths capable of reaching the end-of-day goal. For a time budget B the objective becomes:

$$\max_{D} \; H(D) \quad \text{S.T. } C(D) \leq B \tag{8}$$

This conservatively subordinates the mapping task to navigation goals. It permits a modular integration of corridor exploration into mixed mission plans containing prescripted sampling and navigation waypoints. Despite some limitations, corridor exploration is well-suited as a representative test case for autonomous site survey. It involves significant resource constraints yet remains tractable for real-time implementations.

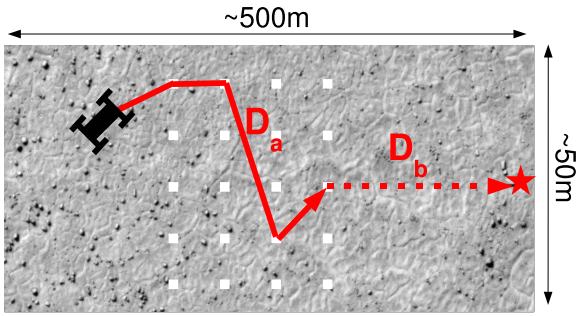

Figure 5. The corridor exploration problem.

3.2. Adaptive Path Selection

Our corridor exploration solution splits the plan into a near-term discretionary portion D_a and a fixed remainder D_b that drives to the goal. The rover begins with the exploration budget allocated evenly among D_b; as it advances it reallocates an appropriate portion of this budget to the time pool available for discretionary planning.

A local planner allocates the near-term discretionary budget $C(D_a)$ by considering several lateral offsets for each forward position along the corridor (Figure 5 right). We search the space of lateral offsets using a recursive greedy path planning algorithm first formulated by Chekuri and Pal [21] and elaborated by Singh et al. [20]. The algorithm selects a midpoint that breaks the path planning problem into two smaller segments and evaluates potential assignment ratios of the remaining resource budget among the halves. For each candidate resource allocation split we recursively call the same routine to identify the best subpath. We evaluate plans by simulating periodic observations along the path and computing their joint entropy with the Gaussian process model. Candidate paths must not exceed the budget allocation for their segment; if there is no valid option the planner defaults to an infinite cost path driving directly to that segment's endpoint.

The short planning horizon permits fast real-time replanning; for a modern single-core laptop processor running unoptimized code the entire replanning procedure requires fewer than 30 seconds. In the physical experiments the rover replanned its path at 3 minute intervals, evaluating paths that each simulated on the order of 50 future observations. Note that the algorithm is myopic; it ignores POMDP aspects of the problem and presumes that model parameters are correct at each timestep.

4. FIELD EXPERIMENTS

Amboy Crater is a basaltic lava flow in the Mojave Desert of California. It is $70km^2$ in area and consists of vesicular pahoehoe lavas deposited over several distinct volcanic events (Figure 1). It is especially interesting to planetary geologists because the flows may be analogous to surfaces on the Moon and Mars [22]. Experiments at Amboy Crater demonstrate our corridor exploration algorithm in kilometer-scale site survey operations. Trials consider the Eastern edge of the Amboy lava flow where eroded basalt flows contact a sediment-covered playa. The two distinct types of surface material — sediment and basalt — correspond roughly to patches of light and dark albedo visible in the overflight image of Figure 2. We performed several trials in which the rover autonomously maps these surficial units.

4.1. Navigation

Zoë (Figure 1) is a field robot developed at Carnegie Mellon University to test mission scenarios involving long-range navigation [23]. It is 2 meters in length and can travel up to 1 meter per second under solar power. A dual passive axle design allows steering by driving the wheels at different velocities to produce variable-radius drive arcs while at the same time accommodating obstacles up to a 20 centimeters in height.

Mast-mounted navigation cameras provide stereo vision for obstacle avoidance and terrain evaluation based on local slope and roughness. Onboard navigation software uses these terrain evaluations along with a D* path planning algorithm [24] to identify optimal paths between prespecified waypoints. Our software architecture treats terrain-based navigation separately from science-driven path planning. Autonomous science software gives each path segment in the form of a single end waypoint, and the terrain-based navigation algorithm attempts to reach a goal area within 5 meters of this location. This does not guarantee the rover will follow the exact path anticipated by the science planner; the rover occasionally detours to avoid large bushes appearing throughout the corridor.

Zoë's localization strategy emulates planetary exploration scenarios. We initialize rover position through differential GPS to simulate manual landmark-based techniques for finding accurate start-of-day positions. During the traverse the rover updates its position estimate with dead reckoning. Planetary rovers can derive absolute heading data from celestial and solar navigation systems utilizing ephemeris calculations; we simulate this capability using periodic heading corrections with an onboard magnetic compass. Some localization error persists due to wheel slip; depending on the terrain this error ranges from 0 to 5% of distance traveled.

4.2. Spectrum Collection and Interpretation

In the Amboy experiments Zoë's onboard science package consists of a Visible Near-Infrared (VIS/NIR) reflectance spectrometer that can collect data from distant solar-illuminated targets. The spectrometer's objective lens, or *foreoptic*, is mounted to a pan-tilt servo unit fixed to the rover mast. A calibration target mounted to the deck provides a white reference to determine targets' absolute reflectivities despite lighting changes. The rover

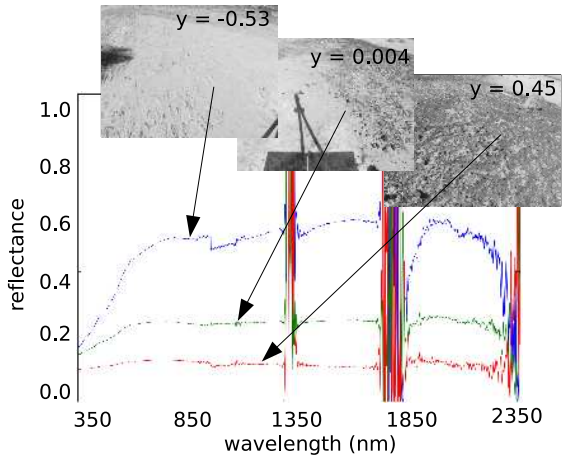

Figure 6. Extreme and mean observations from the traverse of Figure 3 with associated images and spectra. We exclude noisy H_2O absorption bands near 1350*,* 1850*, and* $2350nm$*.*

recalibrates against the white reference every 7 minutes using a scripted sequence described in a companion paper of these proceedings [25]. During normal operation the reflectance spectrometer aims at a fixed -30° angle of declination at the ground directly in front of the rover. At this range its field of view projects to an ellipse measuring approximately 14 and 6 centimeters along the long and short axes. We acquire spectra at 7 second intervals during forward motion, resulting in periodic measurements of surface material from the rover path.

The spectra each contain over 2000 bands, but since they describe a mixture of two physical materials (basalt and sediment) their intrinsic dimensionality is low. A single-output Gaussian process — which only predicts scalar observations — can still model these spectra effectively. Reflectance spectra of macroscopic mixtures are generally linear combinations [26], and linear dimensionality reduction recovers the principal distinction. We compute averages over 4 bandwidth windows to yield a single feature vector, which is rescaled and further reduced to a single dimension using Principal Component Analysis. The end result is the desired scalar observation: the basis derived from a linear projection that permits the best possible reconstruction of collected spectra. This method does not require labeled samples or other *a priori* knowledge about the spectra the robot will encounter. Figure 6 shows mean and extreme observations from the original example traverse along with context imagery.

4.3. Remote Sensing

In addition to the *in situ* spectroscopy Zoë utilizes remote sensing data from one of two sources. The ASTER instrument provides georeferenced VIS/NIR data, yielding 15m per pixel resolution and registration accuracy on the order of 1 pixel. To compensate for registration errors we apply a Gaussian blur with a 1 pixel standard deviation. This low-pass filter preserves coarse information while excluding fine boundary detail that is likely to be inaccurate due to localization error. The ASTER images provide three VIS/NIR channels. Our experiments also investigate a remote sensing data product based on a USGS Digital Orthophoto Quadrangle (DOQ), a single-channel overflight image data product in the visible spectrum [27]. The resolution of DOQ data is 1 meter per pixel, making it analogous to high-resolution instruments such as the HiRISE camera [28]. Registration accuracy is negligible, but we apply the same Gaussian blur strategy in order to compensate for localization drift.

4.4. Procedure

The Amboy field expedition visited three different locales with over 100 traverses ranging from 100m to 1km. Here we focus on a single 300 meter-long exploration corridor. Figure 2 shows DOQ overflight imagery of the traverse area together with a rover plan produced after 27 spectral samples (colored dots). The path planning algorithm chooses a strategy that covers the principal units of surface material within the corridor while respecting the time budget. The blue line indicates one isocontour of the marginal prediction entropy; in this case entropy of the black basalt is high because no sample from this patch has been collected. The entire suite of tests compared several mapping techniques:

- **Adaptive DOQ**: An exploration scheme utilizing the 1m/pixel overflight data in one visible band. We derive class labels for individual map pixels by transforming the final Gaussian process predictions into two classes using k-means quantization ($k = 2$).
- **Adaptive ASTER**: A method using three-band ASTER VIS/NIR orbital data with 15m/pixel resolution. K-means produces final pixel labels.
- **Fixed Transect**: An exploration scheme consisting of a straight drive across the exploration corridor. K-means produces final pixel labels.
- **Fixed Coverage**: An exploration scheme consisting of a zig-zag coverage pattern to uniformly cover the exploration corridor, subject to time constraints. K-means produces pixel labels.
- **Synthetic Random**: A map that does not correspond to any physical traverse; here both pixel class labels are chosen with equal probability.
- **Synthetic Uniform**: Another synthetic map, filled with pixel labels corresponding to the predominant class.

Time budgets were 24 minutes for each physical exploration method except the trials with DOQ overflight data.

Method	Accuracy (std)	Features (std)	p-value (n)
Adaptive			
DOQ	0.87 (0.01)	199.3 (45.3)	<0.01 (3)
ASTER	0.81 (0.03)	225.0 (32.8)	<0.05 (4)
Fixed			
Coverage	0.75 (0.05)	206.0 (14.4)	n/a (4)
Transect	0.74 (0.09)	75.9 (15.0)	0.45 (7)
Synthetic			
Uniform	0.57 (< 0.01)	n/a	<0.01 (5)
Random	0.50 (< 0.01)	n/a	<0.01 (5)

Table 1. Results of the Amboy Crater Traverses: map reconstruction accuracy, the average number of features, and the p-value of a one-tailed T-test against the methods' equivalence to a fixed coverage pattern.

Two of these DOQ trials used a budget of 30 minutes, and a third (the example appearing in Figures 2 and 3) was handicapped with a budget of only 16 minutes. The expedition ended before we could perform time-equivalent DOQ tests, so we encourage caution in interpreting the statistical significance of the DOQ result.

The accuracy score measures the fidelity of a reconstructed map based on the samples collected by each method. We evaluate each dataset by retraining a Gaussian process model using the high resolution remote sensing image. We associate these pixel-wise classifications with their most numerous matches in a hand-labeled copy of the orbital image. This yields an accuracy score for the reconstructed two-class map.

5. RESULTS

The onboard adaptive exploration algorithm adapted to identify informative paths along the exploration corridor. Figure 3 shows the end of the trial after the rover has traveled 0.46 kilometers. In this trial, which lasted approximately 24 minutes, the rover chose a modified coverage pattern that found an appropriate compromise between boundary-following and coverage. This image illustrates the inference result with areas of dense basalt in yellow and clay sediment in green. The map prediction is uncertain in areas that are far from the rover path (consider the unvisited basalt patch in the lower right), but more accurate within the corridor. Samples generated by the path planning algorithm offer a representative sample from which a human could more easily construct a complete interpretation of the environment.

Table 5 shows performance statistics resulting from 18 trials in this traverse area. Reported p-values result from a one-tailed T-test against the hypothesis that performance is equivalent to a static coverage pattern. Figure 7 plots the number of features collected against reconstruction error. In general the adaptive methods outperform static methods, which in turn outperform the uniform and random cases. Traverses based on DOQ overflight data perform best of all. DOQ trials are subject to the timing

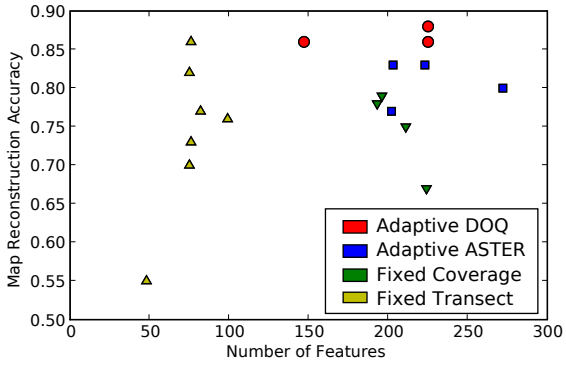

Figure 7. Results of the experimental trials from Amboy Crater: the number of spectrometer samples collected in each run and the accuracy of the reconstructed map.

issues mentioned above, but we include this data here for several reasons. The low variance among the DOQ accuracy implies that the time budgets were not different enough to pose a clear handicap or advantage. Moreover, due to navigation irregularities between runs, the average number of measurements collected was actually *smaller* for DOQ trials than for the coverage patterns. The coverage patterns happened to drive through open terrain, resulting in fewer obstacle-induced pauses.

Several trends are apparent. First, adaptive methods significantly outperform static coverage patterns for $\alpha = 0.05$. The variance in the number of features collected is greater for adaptive methods than for fixed methods, while the accuracy variance is reduced. This discrepancy — greater variation in path choice, and less variation in reconstruction results — might be related to the adaptive methods' responsiveness to navigation obstacles and spectrometer imaging conditions.

6. CONCLUSIONS AND FUTURE WORK

The Amboy Crater experiments demonstrate a system that performs automated site survey on kilometer scales. It begins with virtually no *a priori* knowledge of surface materials. Instead it learns the relevant distinctions on the fly and improves its exploration efficiency by leveraging learned trends in spatial proximity and correlations with remote sensing. The corridor exploration task exhibits some artificial constraints, but demonstrates the essential features of autonomous site survey.

Future work will also consider the selective data return problem from the perspective of Bayesian experimental design. While the execution is subtly different, experimental design serves as a guiding principle for choosing optimal observation subsets. We will also consider more sophisticated feature classification methods; unsupervised dimensionality reduction techniques could also be used to distinguish many geologic classes for multiple-output Gaussian processes. As the space of collected

spectra becomes more diverse, nonlinear dimensionality reduction techniques may prove more appropriate. Finally, we aim to integrate the Gaussian process models discussed here with intelligent spectra collection utilizing rock detection and visual servoing [25]. While there is much scope for development, the Amboy experiments provide a promising proof of concept for kilometer-scale site survey by autonomous spacecraft.

ACKNOWLEDGMENTS

We thank Francisco Calderón, Dom Jonak and James Teza for technical assistance. Ron Greeley, Phil Christensen, and Shelby Cave helped us choose and understand the Amboy field site. NASA ASTEP grant NNG0-4GB66G supported the work.

REFERENCES

[1] R. Castaño, R. C. Anderson, T. Estlin, D. DeCoste, F. Fisher, D. Gains, D. Mazzoni, and M. Judd, "Rover Traverse Science for Increased Mission Science Return," in *IEEE Aerospace*, March 2005.

[2] S. Chien, R. Sherwood, D. Tran, B. Cichy, C. R. G., R. Castaño, A. Davies, D. Mandl, S. Frye, B. Trout, *et al.*, "Using Autonomy Flight Software to Improve Science Return on Earth Observing One," *Journal of Aerospace Computing, Information, and Communication*, vol. 2, pp. 196–216, 2005.

[3] A. Castaño, A. Fukunaga, J. Biesiadecki1, L. Neakrase, P. Whelley, R. Greeley, M. Lemmon, R. Castaño, and S. Chien, "Automatic detection of dust devils and clouds on Mars," 2007.

[4] C. Williams and C. Rasmussen, "Gaussian Processes for Regression," *Advances in Neural Information Processing Systems*, vol. 8, pp. 514–520, 1996.

[5] C. Paciorek and M. Schervish, "Nonstationary Covariance Functions for Gaussian Process Regression," *Advances in Neural Information Processing Systems*, vol. 16, pp. 273–280, 2004.

[6] T. Pfingsten, M. Kuss, and C. E. Rasmussen, "Nonstationary Gaussian Process Regression using a Latent Extension of the Input Space," 2006.

[7] M. Shewry and H. Wynn, "Maximum Entropy Sampling," *Journal of Applied Statistics*, vol. 14, no. 2, 1987.

[8] C. Ko, J. Lee, and M. Queyranne, "An Exact Algorithm for Maximum Entropy Sampling," *Operations Research*, vol. 43, no. 4, pp. 684–691, 1995.

[9] J. Lee, "Constrained maximum-entropy sampling," vol. 46:5, pp. 655–664, Sep. - Oct. 1998.

[10] C. Guestrin, A. Krause, and A. Singh, "Near-optimal Sensor Placements in Gaussian Processes," School of Computer Science, Carnegie Mellon University, 2005.

[11] T. Loredo, "Bayesian adaptive exploration in a nutshell," *PHYSTAT*, 2003.

[12] R. Simmons, D. Apfelbaum, W. Burgard, M. Fox, D. an Moors, S. Thrun, and H. Younes, "Coordination for multi-robot exploration and mapping," *AAAI*, 2000.

[13] L. Pedersen, M. D. Wagner, D. Apostolopoulos, and W. R. L. Whittaker, "Autonomous robotic meteorite identification in antarctica," in *ICRA*, May 2001.

[14] S. Moorehead, R. Simmons, and W. Whittaker, "A multiple information source planner for autonomous planetary exploration," in *ISAIRAS*, 2001.

[15] D. MacKay, "Information-Based Objective Functions for Active Data Selection," *Neural Computation*, vol. 4, no. 4, pp. 590–604, 1992.

[16] V. V. Fedorov, *Theory of Optimal Experiments*. academic press, 1972.

[17] P. Sebastiani and H. P. Wynn, "Maximum Entropy Sampling and Optimal Bayesian Experimental Design," in *Journal of the Royal Statistical Society*, 2000.

[18] Trey Smith and David R. Thompson and David Wettergreen, "Generating exponentially smaller pomdp models using conditionally irrelevant variable abstraction," *ICAPS*, 2007.

[19] A. Meliou, A. Krause, C. Guestrin, and J. M. Hellerstein, "Nonmyopic informative path planning in spatio-temporal models," Tech. Rep. UCB/EECS-2007-44, UC Berkeley.

[20] A. Singh, A. Krause, C. Guestrin, W. Kaiser, and M. Batlin, "Efficient Planning of Informative Paths for Multiple Robots," in *IJCAI*, 2007.

[21] C. Chekuri and M. Pal, "A Recursive Greedy Algorithm for Walks on Directed Graphs," in *IEEE Symposium on the Foundations of Computer Science*, 2005.

[22] R. Greeley and J. D. Iversen, "Field Guide to Amboy Lava Flow, San Bernardino County, California," 1978.

[23] D. Wettergreen et al., "Second experiments in the robotic investigation of life in the atacama desert of chile," in *ISAIRAS*, 9 2005.

[24] S. Singh, R. Simmons, T. Smith, A. Stentz, V. Verma, A. Yahja, and K. Schwehr, "Recent Progress in Local and Global Traversability for Planetary Rovers," in *ICRA*, 2000.

[25] F. C. Peralta, D. R. Thompson, and D. Wettergreen, "Autonomous Rover Reflectance Spectroscopy with Dozens of Targets," 2008.

[26] Mustard, J.F. and C. M. Pieters and S. F. Pratt, "Deconvolution of spectra for intimate mixtures," in *LPSC XVII*, 1986.

[27] United States Geological Survey, *Digital Orthophoto Quadrangles Fact Sheet 057-01*.

[28] A. S. McEwen, W. A. Delamere, E. M. Eliason, J. A. Grant, V. C. Gulick, C. J. Hansen, K. E. Herkenhoff, L. Keszthelyi, R. L. Kirk, M. T. Mellon, S. W. Squyres, N. Thomas, and C. Weitz, "Hirise: The high resolution imaging science experiment for mars reconnaissance orbiter," in *Lunar and Planetary Institute Conference Abstracts*, vol. 33, p. 1163, Mar 2002.