

Efficiently Sampling from Underlying Models

Greydon Foil, David Wettergreen

Robotics Institute

Carnegie Mellon University

Pittsburgh, PA 15213

Email: gfoil@alumni.cmu.edu, dsw@ri.cmu.edu

Abstract— The capability and mobility of exploration robots is increasing rapidly, yet missions will always be constrained by one main resource: time. Time limits the number of samples a robot can collect, sites it can analyze, and the availability of human oversight, so it is imperative the robot is able to make intelligent actions when it comes to choosing when, where, and what to sample, a process known as adaptive sampling.

This work advances the state of the art in adaptive sampling for exploration robotics. We take advantage of the fact that rover operations are typically not performed in a vacuum; extensive contextual data is often present, most often in the form of orbital imagery, rover navigation images, and prior instrument measurements. Using this context, we apply Bayesian and nonparametric models to decide where best to sample under a limited budget, using real X-ray lithochemistry data.

We find that our methods improve both the diversity of samples collected as well as select samples that are representative of the dataset. We find that model-based approaches made scalable with Dirichlet processes improve sampling results when the underlying number classes and class distribution is unknown. Unlike previous works, our approaches reduce the impact of noise on sampling location, a common problem when selecting samples based on noisy or incomplete contextual data.

I. INTRODUCTION

As exploration robots increase in capability and mobility, missions are measured less by what a robot can do and more by how, why, and where a robot can make the most progress towards its mission objectives. Time is the constraining resource and so we seek to increase productivity.

Adaptive sampling can improve the utilization of time. Adaptive sampling refers to techniques that use onboard models and statistical analysis to choose where and how to deploy instruments and acquire measurements. Adaptive sampling methods can be used to detect anomalous features, plan paths with stops at diverse types of terrain, or analyze an interesting sample and decide to take more detailed measurements. These techniques can increase data acquisition rates and quality, improving the overall yield of a mission.

Current Mars rover operations give an example of the need for adaptive sampling techniques. Daily plans budget specific operations down to the minute, and each operation is given a fixed duration, sometimes only a few hours. However, the capabilities of the rover far exceed the limited plans. For example, the Mars 2020 mission is testing the Planetary Instrument for X-Ray Lithochemistry (PIXL), a spectrometer used to analyze rocks on a microscopic scale. A quick scan of a postage stamp-sized surface can take an hour but is very noisy, yet reducing this noise with a full

scan of the surface takes multiple days. Scheduling constraints and communication delay prohibits scientists from reviewing data in real-time and they are thus unable to follow up at any promising sample locations. However, adaptive sampling techniques enable the rover to quickly and autonomously choose a subset of interesting points for detailed follow-up.

In order to make the most of robotic capabilities, it is necessary to develop adaptive sampling techniques that optimize the quality of the sampling data collected by robots while simultaneously allowing for greater exploration. Robots have improved at detecting objects [8, 21, 31, 12], environmental phenomena [6], or locating promising survey locations only using low-resolution orbital imagery [32]. However, these methods are incomplete or only execute under certain conditions, causing the rover to miss interesting objects or measurement locations during a traverse, like those cases shown in Figure 1. The frustration of the science team to have gotten so close and have just missed can hardly be overstated.

The objective of this work is to introduce adaptive sampling methods that increase the diversity of samples collected and collect samples that are representative of the scene as a whole. We accomplish this goal by making use of rover measurements that can be gathered quickly or with little effort, such as images or spectral scans, to estimate models of the domain, then sample from these models in order to confirm expected measurements and spot interesting areas that deviate from known models.

We introduce multiple methods for improving adaptive sampling when provided with contextual scene information or when no information is provided. Our first methods are model-based, and we find that using Dirichlet process priors to grow the models over time improves our sampling results. Our final methods are based on Bayesian optimization, building approximations of the underlying reward function at unsampled locations which are robust to noise in the contextual data.

This work maximizes two objectives in an *efficient* manner: sample *diversity* and collecting *representative* samples. Here efficiency is defined as the amount of information gained per unit of time. This is maximized by focusing observation on areas expected to improve our knowledge and not oversampling understood areas. We define diversity as a measure of the different types of samples collected, aiming to collect the most widespread set of samples possible. Representative samples are those that are characteristic of the sample space. These samples provide the most information about the composition



Fig. 1. Two instances of where better target selection could have improved robotic sample selection. In both situations the robot navigates to a location and collects a grid of spectra immediately in front of it. In the left image it has correctly identified a region of interest and navigated there, but is pointing in an incorrect direction and samples from a mixture of the materials. In the right image it has again navigated to the correct location but aims its spectrometer beyond the optimal targets, including patches of white salts that are of interest to the science team.

sample space.

II. RELATED WORK

Adaptive sampling is a broad topic that shares common themes with information theory, active learning, and Bayesian optimization. Information-theoretic approaches have been used in robotic sampling tasks, and their methods aim to learn information about parameters in environmental models by choosing sample points that reduce the posterior model uncertainty.

Information gain was used by Pedersen et al. [21] to decide whether to deploy instruments to analyze a prospective meteorite. Shewry and Wynn introduce the idea of Maximum Entropy Sampling (MES) [28, 5, 25] to find the best set of sampling points within the collection of all points that have the maximum joint entropy. Krause et al. [14] propose an alternative using a mutual information criterion in their work on finding optimal sensor placements within a map. Mutual information shares similarities with MES but removes any reward from the observations themselves, causing the metric to avoid high-entropy locations on the edge of the map.

Active learning is closely related to adaptive sampling. In active learning, data is presented to an agent, then the agent is allowed to query for more information on a single data point.

There are two main methods used in active learning. The first is *pool-based sampling* [15], in which all of the data is visible to the agent, and the agent must select a sample point to query within that full set. This is the case most commonly considered in active learning research, and is similar to the sampling with context experiments in Section V-A in which the full data is observable and we must select the point within this data to sample.

The second is *stream-based sampling* [1, 4], in which a single observation is available at a time, and the agent must decide whether to query that point or discard it. This is the case in the limited context experiments in Section V-B, in which we model a microscopic raster of a surface in which the instrument must decide immediately whether or not to fully scan a location.

Various active learning approaches quantify the informativeness of unlabeled data using metrics such as model uncertainty [15], expected model change [27], error reduction [23], or reduction in model variance [4]. A good overview of active learning is provided in [26].

Gaussian processes and Bayesian optimization techniques have also been used extensively for adaptive sampling in the spatial modeling domain. Thompson uses a Gaussian process with orbital maps as a latent input to model the expected sample readings at future sites, demonstrating improvements in autonomous rover sample collection [31]. Foil et al. use Gaussian processes to model subsurface hydrogen, as measured by a rover-mounted neutron spectrometer, providing scientists more accurate maps to guide rover navigation [9].

We note the difference between spatial design and *sequential design* [34, 3]. In spatial design, the experiment is designed before samples are collected. In our work we instead focus on sequential design, in which the process is iterative, and collected samples alter future rewards.

We note that this problem shares similarities with next-best view optimization in computer graphics and robotics literature. In that situation, a model of some object or environment is desired, and a robot or agent must choose the camera or sensor views that provide the greatest increase in information about the object. This estimate of information content is typically done using information gain, entropy, or other statistical measures similar to our own. A survey of approaches is provided in [24].

Haines and Tao Xiang [11] considers a pool-based active learning case which aims to jointly discover rare classes from within a dataset and refine the parameters of a classifier over the unsampled data, by the use of Dirichlet processes as a nonparametric estimator on the probability of any sample coming from an unknown and, as yet, unobserved class. The goal then is to balance finding new classes with refining existing ones. Many approaches, like Pedersen et al. [21], find points that have high uncertainty with regards to the existing models, but most metrics, such as entropy, do not account for




Fig. 2. Example data from the Planetary Instrument for X-Ray Lithochemistry instrument. Left: A microscopic image of the surface of the “Troughite” sample. Center: A false-color image showing the responses of the elements Al (red), Ca (green), and Ti (blue). Right: A classification of the surface by clustering the spectra at each point. From [33].

unknown classes. Other methods, such as sampling points with high probability of misclassification with regards to known class models, have done well at refining models, but do not account for unknown models without explicit consideration.

The overall goals of Haines et al.’s work mirror our own, however, while they consider a classification approach, in which each point of the data has a pre-assigned class label, our work has no such labels, and must instead estimate class distributions only from prior observations.

III. MOTIVATIONS AND DATASET

We evaluate our methods with real data collected by the Planetary Instrument for X-Ray Lithochemistry (PIXL). PIXL is an X-ray fluorescence spectrometer that collects a raster of a microscopic surface, an example of which is shown in Figure 2. It can collect a noisy spectrum in seconds, or integrate measurements for minutes to collect a longer, more accurate spectrum. Due to mechanical constraints, PIXL is unable to return to a specific point with a high enough accuracy to guarantee the location is in the instrument’s best field of view, so it must collect a quick spectrum and immediately decide whether or not the content of that measurement warrants further integration to improve the signal-to-noise ratio.

If each follow-up action was vetted by scientists on Earth, a single raster would take days or months to complete. Instead, PIXL requires adaptive sampling methods, allowing the rover to analyze measurements as they are collected and determine the appropriate follow-up operation. This greatly reduces the need for scientist oversight, as well as builds environmental models over time that can be used at future sample sites to reduce the number of redundant samples collected.

The data used in this work was collected using the PIXL instrument at the Jet Propulsion Laboratory. The sample used here is from the Dressermats dataset, a sample specimen analyzed as a raster of 180 by 170 points, in which each point is a 4096-element spectrum. Each spectrum was collected using a full-duration scan, meaning the instrument integrated data over a longer period of time in order to increase the

signal-to-noise ratio. The data is integrated over specific sets of wavelength ranges corresponding to certain elements, such as Al, Si, and Ti. The end result is a 9-element vector at each measurement location that approximates the abundance of core elements at that location. For more information on the data and pre-processing routine, see [10].

Collecting a dataset at this resolution and quality is a time-consuming process; individual measurements by the device take anywhere from 15 to 60 seconds, so a quick raster of a surface takes an hour or two, while the more accurate raster of long-dwell spectra takes many tens of hours. It is thus infeasible to do full-duration scans at all locations, as there are severe time constraints on sampling on planetary missions. Instead, adaptive sampling methods can be used to analyze noisy samples, which can be taken quickly, then choose a subset of the most promising points at which to collect long-dwell spectrum.

IV. APPROACH

Our adaptive sampling scenario requires methods that have three main qualities: they take advantage of underlying trends in the data, they can distinguish noise from meaningful changes in physical variables, and they are flexible in sampling from an unknown number of underlying classes.

Two methods satisfy these characteristics. First, we introduce a mixture model-based approach and its extension to use a Dirichlet process prior to generate new classes and predict misclassifications. These mixtures allow the model to characterize different mineralogical classes, estimate the noise present in each class, and flexibly expand to create room for new classes as evidenced by the data.

Second, we introduce a method that estimates the reward of future samples using Gaussian processes (GPs), allowing us to use established Bayesian optimization techniques to select points that have a high expected reward. This approach is less susceptible to noise in the model, due to smoothing reward estimation over points with similar features, while also providing an actionable measure of model uncertainty.

A. Adaptive Gaussian Mixture Models

Our first method is based on the work of Sillito et al. [29] and their incremental kernel density estimation. We consider the creation of a mixture model, similar to a Gaussian mixture model (GMM) [19].

Sillito et al. consider the specific case in which a GMM is incrementally generated and cannot be learned a priori using training data. They propose to have a limited number of mixtures, initialized with very loose parameters. As data is collected, the mixtures are gradually adjusted to include new data and similar mixtures are merged, such that the final result resembles a GMM.

Due to our problem formulation, in which points are selected incrementally and contextual data is only a proxy for the real information at a point and not necessarily representative of the data at that point, an approach such as this is flexible enough to allow a model-based formulation, while also being lightweight enough to work with a very limited number of points. Because the covariances of each mixture are initialized to be very loose and are not based on the values of the points themselves, the model can grow in a stable manner, unaffected by mixtures with very constrained covariances.

This approach is referenced as **AGMM** in our tests. Its formulation is computationally cheap, it is flexible enough to work well at finding and sampling diverse points in the data, and it works well with a very limited amount of training data. However, like conventional GMMs, it still assumes prior knowledge of the appropriate number of mixtures in the model. Too few mixtures leads to an over-generalization of the data, while too many mixtures can lead to over-fitting.

B. Adaptively Adding Classes Using Dirichlet Process Priors

We propose a method for adaptively increasing the number of mixtures in our model by using a Dirichlet process prior [7]. Instead of estimating the number of classes in the AGMM model, we use a Dirichlet process prior to estimate whether or not a newly-collected sample point should be added to an existing class or added to an entirely new class.

Our formulation assumes a similar structure to that of the Chinese Restaurant Process (CRP). In that scenario, patrons enter a Chinese restaurant one by one. Inside the restaurant are infinitely many tables, $T = \{t_1, \dots\}$, and a customer chooses a table with probability:

$$x_n | x_1, \dots, x_{n-1} = \begin{cases} t_k & \text{with probability } \frac{num_{t_k}}{n-1+\alpha} \\ t_{new} & \text{with probability } \frac{\alpha}{n-1+\alpha} \end{cases} \quad (1)$$

with classes $k \in K_+$, Dirichlet concentration parameter α , and num_{t_k} is the number of patrons sitting at table t_k .

In this case, all tables and all patrons have equal probability, and thus the probability of being at a new table or an old table are only dependent on the number of patrons already at the restaurant and sitting at the table, respectively.

However, we can leverage the mixture model probabilities in order to obtain a more accurate distribution of our data. To demonstrate this, we introduce the *Food Court Process*.

Imagine that instead of a restaurant with up to an infinite number of tables, there instead is a food court with up to an infinite number of restaurants. Patrons enter the food court, then choose a restaurant at which to eat.

In the CRP example, patrons are equally probable and their choice of table is based entirely on the Dirichlet concentration parameter, α , and the number of patrons at each table. In the food court case, imagine instead the more realistic scenario where each patron, y , has a preference of food, and individual groups of people are more or less likely to appear. In this case, the model can be updated to look like this:

$$x_n | x_1, \dots, x_{n-1} = \begin{cases} r_k & \text{w/ prob. } \frac{num_{r_k} P_{r_k}(y | r_k)}{n-1+\alpha} \\ r_{new} & \text{w/ prob. } \frac{\alpha}{n-1+\alpha} P(y) \end{cases} \quad (2)$$

in which the probability of a patron choosing an existing restaurant is proportional to the popularity of the restaurant, weighted by their preference to that type of food. If the patron does not like any of the restaurants that already have patrons, then it is more likely they will choose to eat at a new one.

This is the basis for our **AGMM-AddClasses** approach. As we sample points, our mixture model builds up a number of distributions based on the data in the sample points belonging to those mixtures. When a new sample point arrives, the algorithm must decide if it should add a point to an existing mixture or create a new one, using the aforementioned equation:

$$P(c \in \{C, new\} | x) \propto \begin{cases} \frac{m_c}{n-1+\alpha} P_c(x | c) & \text{if } c \in C \\ \frac{\alpha}{n-1+\alpha} P(x) & \text{if } c = new \end{cases} \quad (3)$$

for sample x , the set of mixtures C , and m_c being the number of points in mixture c . $P(x)$ is the prior probability of the sample and $P_c(x | c)$ is the probability of our sample belonging to mixture c , given by our mixture model.

Intuitively, one might imagine that, because we are sampling the most diverse set of points we can find, we will always be likely to create new mixtures, as sampled points are selected based on their “distance” from existing mixtures. In reality, however, due to the noise in the system this is not always the case. Samples that look promising given the contextual data may only appear promising because of a large amount of noise in that samples’s measurement. After fully observing the point it may be the case that that sample is not novel enough to warrant the creation of a new mixture.

C. Probability of Misclassification

Finally, we augment the AGMM formulation with an approach similar to that of Haines et al. [11]. In this work they are interested in calculating the probability of any point being incorrectly classified given an existing model. This has been considered in prior work, such as Lewis and Gale’s work on text classification [15], in which they propose selecting sample points with class probabilities closest to 0.5. However, their work doesn’t consider the probability of a point belonging to an unobserved case.

Haines et al. instead considers what they describe as $P(Wrong)$, or the probability of being incorrectly classified

given all existing classes, as well as the underlying probability that it belongs to a new class:

$$P(Wrong|x) = 1 - P_n(c'|x) \quad (4)$$

$$c' = \arg \max_{c \in C} P_c(x|c) \quad (5)$$

where P_n is the Dirichlet-weighted probability of belonging to an existing class or a new, unobserved class:

$$P_n(c \in \{C \cup new\}|x) \propto \begin{cases} \frac{m_c}{n-1+\alpha} P_c(x|c) & \text{if } c \in C \\ \frac{\alpha}{n-1+\alpha} P(x) & \text{if } c = new \end{cases} \quad (6)$$

We apply a modification of this technique to our sample selection process, referring to it as **AGMM-P(Wrong)**. We consider the case where the mixture model is defined as the AGMM model discussed above. When points are considered for sampling, existing mixtures are used to calculate the probability of misclassification, weighted by the Dirichlet probability of the point belonging to a new class. Similar to Haines, we select the point that maximizes $P(Wrong)$ and add it to our AGMM model.

We hypothesize that this approach will provide high overall coverage and sample diversity. By selecting points which are poorly represented by the existing model, the set of selected samples is more likely to contain dissimilar samples that span the search space. This selection method also leverages the Dirichlet process by choosing points that are likely to belong to a new, undiscovered class.

D. One-Step Bayesian Optimization

In addition to our mixture model work, we introduce an alternative sample selection process that draws from Bayesian optimization approaches to find new sample selection points.

Bayesian optimization is a field of research focused on using Bayesian methods to model reward functions, then sampling methodically within those models to maximize reward. Bayesian optimization is applicable when the direct function to optimize is difficult to model or unknown, or evaluations are expensive, but when samples can be collected at given values in order to estimate the structure of the underlying function. The underlying function is modeled by a known *surrogate function* or *response surface* that is computationally easy to evaluate and analyze, typically represented by a Gaussian process model [22]. As data is collected the Gaussian process is recalculated and the mean and covariance functions adjust to fit to the new data, providing greater accuracy in regions of the search space that have been sampled.

Sampling within Bayesian optimization is a trade-off between exploration, in this case sampling from the unknown regions of the response surface, and exploitation, sampling regions of the response surface known to be promising. The choice of where to sample is typically handled by an *acquisition function* which provides an estimate of the “reward” of a given sample. Optimizing this function typically is done using the principles of maximum expected utility, or minimum expected risk [2].

Our method uses Bayesian optimization to estimate the differential entropy reward of sampling at any point in our contextual data. We hypothesize that the Gaussian process will act as a smoothing operation over similar points, providing a more accurate estimate of reward that is more resistant to noise in the data. Additionally, the variance estimates provide a measure of uncertainty that allow us to consider explicitly the trade off between true exploration or exploitation.

It is important to note that Bayesian optimization typically operates on a single, static objective function. This function does not change over time or with observations, and, thus, repeated queries into the function build upon prior knowledge and can be used to improve the accuracy of the response surface. This is often not the case in our test scenario, as our objective function is the potential reward from collecting a given sample. The information gained in sampling affects that reward, necessitating a re-estimation of the response surface. Thus, our acquisition functions and the parameters chosen for our response surface focus on reliably maximizing the reward of a single sample.

We call this *One-Step Bayesian Optimization*. We consider the problem of selecting the next sample with the highest expected entropy, x_{n+1} , given our sample set X :

$$x_{n+1} = \arg \max_{x_i} H(X \cup x_i), \quad \forall x_i \notin X \quad (7)$$

However, the calculation of entropy is intrinsically dependent on the set of samples, This means that our target objective function is changing at each sampling step. Thus, our objective is to find acquisition functions that provide the greatest return over a single step.

This approach is myopic in nature, only considering the expected reward of the next sample. Alternatively, a number of non-myopic approaches to Bayesian Optimization have been proposed [18, 16] which model the expected reward over multiple samples up to some horizon or cut-off point. In our extended work, which looks at selecting an optimal sampling path for a rover traversal, we consider non-myopic applications of these algorithms to calculate an expected reward with regards to a sampling and time budget [10].

Given an accurate response surface, the Bayesian optimization problem then becomes a selection of an acquisition function. We consider two main acquisition functions, although our work in [10] evaluates others.

First, we consider Expectation of Improvement (EI), in which we calculate the likelihood that we improve over some *incumbent* value, $f(x^+)$, the highest value observed in our objective function thus far [20, 13]. Typically, when multiple samples are collected from the same objective function, this is considered $x^+ = \arg \max_{x_i \in X} f(x_i)$. In our case, where the objective function changes at each iteration, we instead consider x^+ to be the unsampled point that has the highest value of our current model’s mean function: $x^+ = \arg \max_{x_i \in X_{unsampled}} \mu(x_i)$.

EI offers a balanced approach between exploration and exploitation. If the variance of a candidate point is large, it




Fig. 3. An example of the sampling process with context. If context is available, the reward is estimated across all points in the scene.

explores and takes a risk with the hope of finding something rewarding, but if the mean of a candidate point is large, it exploits that information. If we instead want to encode a trade-off between these behaviors, we can include a trade-off parameter, ξ , suggested by Lizotte [17]:

$$EI(x) = \begin{cases} (\mu(x) - f(x^+) - \xi)\Phi(Z) + \sigma(x)\phi(Z) & \text{if } \sigma > 0 \\ 0 & \text{if } \sigma = 0 \end{cases} \quad (8)$$

$$Z = \begin{cases} \frac{\mu(x) - f(x^+) - \xi}{\sigma(x)} & \text{if } \sigma > 0 \\ 0 & \text{if } \sigma = 0 \end{cases} \quad (9)$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ indicate the CDF and PDF of the standard normal distribution, respectively. Lizotte finds good performance with a value of $\xi = 0.01$ and finds that a decaying schedule, as suggested in the PI literature, actually underperforms when compared against using a static value for ξ . We refer to this approach as **GP-EI**.

Note that because $f(x^+)$ comes from our unsampled data, it is possible that it is the point selected by the acquisition function. This is acceptable, as it is intended only to provide a baseline against which to compare the other unsampled points. It is still likely that other sampling locations are selected that have both a high expected value and high variance.

Finally, we evaluate a simple model, **GP-MaxMean**, which ignores sample variance and instead chooses the sample with the maximum expected value in our Gaussian process:

$$x_{n+1} = \arg \max_{x_i \in X_{\text{unsampled}}} \mu(x_i) \quad (10)$$

This is largely a measure of the ability of the Gaussian process to estimate the reward. Additionally, it acts as a smoothing function to reduce the impact of noise in sample-dense regions of our feature space.

V. RESULTS

We conduct two experiments that mimic common adaptive sampling scenarios seen in robotic exploration. The first is *Adaptive Sampling with Context*, in which the robot has contextual information at all possible measurement locations. This is similar in concept to a rover collecting a panoramic image or a quick raster with a spectrometer, then choosing targets of interest for follow-up measurements. Because the full contextual data is available, this experiment is designed




Fig. 4. An example of the sampling process without context. If context is limited, there is no information about upcoming points and the reward is only estimated on prior observed data.

to demonstrate and compare the effectiveness of approaches under a minimal number of sampling constraints.

The second experiment is *Adaptive Sampling with Limited Context*. This experiment considers how to handle sampling in a scenario like that seen by the PIXL instrument, in which a raster of data is collected and evaluated and the robot must decide immediately whether to collect a detailed spectrum [30]. The contextual data is limited to what the rover has seen thus far, mostly consisting of noisy spectra and any more accurate readings from detailed measurements.

Results are scored using two main, complementary metrics: entropy and reconstruction error. Entropy quantifies the amount of diversity in the samples, while reconstruction error quantifies the “coverage” of the sample set.

The entropy score is calculated using differential entropy:

$$H(x) = \frac{1}{2} \sum_{b \in B} \log(2\pi\sigma_b) \quad (11)$$


where $b \in B$ are each of the channels in the data. This score rewards samples that increase the overall information of the sample set. While the channels in the raw spectral data cannot be considered independent, the pre-processing of the data reduces it to nine channels, each measuring elemental abundance, which we can effectively consider independent here.

In contrast, reconstruction error favors sampling points that are representative of the entire dataset. It is calculated as the average non-negative least-squares error when using the collected samples to reconstruct all unobserved samples. Given a set of collected samples $X = \{x_1, \dots, x_n\}$ and unsampled points $Y = \{y_1, \dots, y_m\}$, the error is calculated using the set of weights $W = \{w_1, \dots, w_n\}$ that minimizes:


$$W^* = \arg \min_W \|W^T X - Y\|_2 \quad (12)$$

$$w_i > 0, \quad \forall w_i \in W \quad (13)$$

All results are averaged over a set of 100 datasets randomly-generated from the original PIXL data. Noisy contextual data is generated using Poisson statistics from the original dataset. Each method has a total sampling budget of 20 full-duration samples.



(a) Entropy (Higher is better)



(b) Reconstruction Error (lower is better)

Fig. 5. Results for Adaptive Sampling with Full Context. GP-EI achieves the highest scores under the entropy metric, while the AGMM methods do best with regards to reconstruction error.

A. Adaptive Sampling with Context

The first experiment considers the case when the robot has access to contextual information at all possible measurement locations. This contextual data may be noisy but can be used to approximate a model of the data or model parameters. For example, this data can be used to estimate the Dirichlet concentration parameter, α , in the **AGMM-AddClasses** approach, or the response surface in the Gaussian process-based methods.

Our methods are compared to three established techniques: random sampling, grid-based sampling, and a greedy entropy maximization method, shown in tests as **Max Entropy**. This entropy-based method calculates the expected differential entropy for each point, given the existing sample set. This method does well if there is low noise in the scene, but can greedily sample from noisy locations in the contextual data that only appear rewarding because of noise.

The results for this experiment are shown in Figure 5 and Table I. The AGMM methods all perform competitively compared against the baseline methods. Both **AGMM** and **AGMM-AddClasses** outperform all baseline approaches in both entropy and reconstruction error. **AGMM-AddClasses** demonstrates that there is a strict improvement over the **AGMM** method when classes are added dynamically using a Dirichlet process prior. Keeping the number of classes flexible results in both an increase in entropy and a reduction in reconstruction error.

All Bayesian optimization methods outperform the baseline methods with regards to the entropy scoring metric, indicating that they are well-equipped to find diverse points within the sample data. **GP-EI** achieves the highest overall entropy score out of all tested methods. However, **GP-EI** does poorly when considered under the reconstruction error metric.

In contrast, the **GP-MaxMean** method does well on all metrics. When scored using the entropy metric it outperforms the baseline on both the full- and reduced-channel data,

indicating that using a Gaussian process for nothing other than a smoothing operation has some benefit. When scored using reconstruction error it outperforms all other methods. This is potentially because the smoothing operation leads it to sample fewer outlier points and instead sample points more representative of the overall data.

| Method Name | Max Entropy | Final Entropy | Recon Error |
|----------------------|-------------------------------------|-------------------------------------|-----------------------------------|
| Grid | 3402.2 | 3274.9 | 152.9 |
| Random | 3367.5 \pm 82.9 | 3267.9 \pm 54.7 | 174 \pm 35.8 |
| Maximum Entropy | 3590.9 \pm 76.9 | 3476 \pm 63.4 | 171.5 \pm 22.9 |
| GP-EI | 3728.6 \pm 48.1 | 3636.4 \pm 37.9 | 236.6 \pm 19.8 |
| GP-MaxMean | 3667.8 \pm 41.9 | 3532.1 \pm 46.1 | 152.7 \pm 7.0 |
| AGMM | 3620.7 \pm 61.7 | 3506.9 \pm 49.3 | 159.4 \pm 14.2 |
| AGMM-AddClass | 3667.7 \pm 72.4 | 3545 \pm 45.3 | 152.7 \pm 7.1 |


TABLE I
ADAPTIVE SAMPLING WITH CONTEXT. OUR CONTRIBUTIONS AND THE BEST RESULTS ARE LISTED IN BOLD.

B. Adaptive Sampling with Limited Context

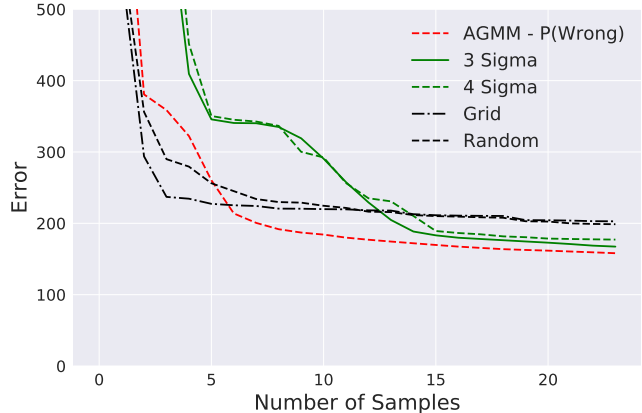
Our second experiment considers the case where the only contextual information available are the points which have already been sampled. Unfortunately, this greatly limits the effectiveness of the model-based approaches used in the full-context experiments.

Instead, we consider the AGMM - P(Wrong) formulation. In this model, there are still a fixed number of K mixtures. However, we are a little more restrictive with the points that we sample, and instead favor only sampling points that are dissimilar from the data we have already collected. At every candidate sample point, we calculate the value of $P(wrong|x)$, as described above, then decide whether or not to sample at that point. For all tests the sampling criterion is

$$\text{trigger when } P(wrong|x) > \tau \quad (14)$$



(a) Entropy (Higher is better)



(b) Reconstruction Error (lower is better)

Fig. 6. Results for Adaptive Sampling with Limited Context. Our AGMM-P(Wrong) achieves very strong results in both entropy maximization and minimizing reconstruction error. The 3 Sigma and 4 Sigma approaches both have a slower ramp-up time, but do well with regards to both metrics.

In this experiment $\tau = 0.997$, chosen to be similar to the interval-based approach with an interval of $[-3\sigma, 3\sigma]$.

We compare to the interval-based approaches proposed by Thompson et al. [33], in which samples are selected if they fall either three or four standard deviations away from the mean sample observed thus far, in any elemental channel. We refer to these methods as **Sigma3** and **Sigma4**, respectively. Additionally, we compare against random- and grid-based selection methods.

Results are shown in Figure 6 and Table II. The overall best performer is the **AGMM-P(Wrong)** method. It achieves the best scores in all but one metric, and even then is close to the optimal score.

Random and **Grid** do very well in this experiment, especially in their initial few samples. This is because both the random and periodic methods sample a larger space of the data, while other methods typically exhaust their sampling budget early in the raster.

The other methods have somewhat of a slower start, as they must build up initial models over the section of the raster where the initialization point is located, so their long-term results are higher, but their initial scores are lower. This is especially true for the interval-based approaches, **3Sigma** and **4Sigma**, as they have to gradually build up a model.

| Method Name | Max Entropy | Final Entropy | Recon Error |
|------------------------|-------------------------------------|-------------------------------------|-----------------------------------|
| Grid | 3347.4 \pm 0.0 | 3278.4 \pm 0.0 | 204.5 \pm 0.0 |
| Random | 3357.0 \pm 43.5 | 3286.8 \pm 34.6 | 202.9 \pm 23.5 |
| 3 Sigma | 3253.2 \pm 47.1 | 3253.2 \pm 47.1 | 174.6 \pm 4.0 |
| 4 Sigma | 3252.4 \pm 34.4 | 3252.4 \pm 34.4 | 180.6 \pm 14.4 |
| AGMM - P(Wrong) | 3352.3 \pm 64.0 | 3332.8 \pm 50.7 | 162.7 \pm 3.9 |

TABLE II

ADAPTIVE SAMPLING WITH LIMITED CONTEXT. OUR CONTRIBUTIONS AND THE BEST RESULTS ARE LISTED IN BOLD.

VI. CONCLUSION

This work has demonstrated a number of adaptive sampling methods that improve a robot’s ability to collect both diverse and representative samples. We have shown that these methods work well in situations where robots have some prior information about a sampling location, as well as in cases where little information is known.

We have introduced three main methods. First, we showed a method for adaptively adding classes to model-based approaches using Dirichlet processes. This approach improved over static models, especially when looking for diverse points in the data.

Second, we demonstrate a method for selecting sample points based on their probability of being misclassified under the current model. This method excels at adaptive sampling with no context, where it is much more difficult to build up a representative model over time.

Finally, we demonstrate the effectiveness of one-step Bayesian optimization, in which a Gaussian process is used to estimate the reward of unsampled points. The Gaussian process helps smooth the reward function over similar points, reducing the impact of noise on the reward estimate, as well as providing an estimate of variance in the reward that can be used to search for diverse points in the data. These methods excelled at finding diverse points.

Our work will continue by applying these methods onboard a planetary rover prototype and conducting field experiments to evaluate the efficiency and effectiveness of adaptive sampling for science exploration.

REFERENCES

- [1] L. Atlas, D. Cohn, R. Ladner, M. El-Sharkawi, R. Marks II, M. Aggoune, and D. Park. Training connectionist networks with queries and selective sampling. *In Advances in Neural Information Processing Systems (NIPS)*, 3:566573, 1990.

- [2] Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.
- [3] D. Busby. Hierarchical adaptive experimental design for gaussian process emulators. *Reliability Engineering and System Safety*, 94(7):11831193, 2009.
- [4] David A Cohn. Neural network exploration using optimal experiment design. *Advances in Neural Information Processing Systems*, 1994.
- [5] N. A Cressie. *Statistics for Spatial Data*. Wiley and Sons, New York, 1991.
- [6] et al. Estlin, Tara A. AEGIS Automated Science Targeting for the MER Opportunity Rover. *ACM Transactions on Intelligent Systems and Technology*, 3:1–19, 2012. ISSN 21576904. doi: 10.1145/2168752.2168764.
- [7] Thomas S Ferguson. A Bayesian analysis of some nonparametric problems. *The annals of statistics*, pages 209–230, 1973.
- [8] Greydon Foil, David R. Thompson, William Abbey, and David S. Wettergreen. Probabilistic surface classification for rover instrument targeting. *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 775–782, 2013.
- [9] Greydon Foil, Terry Fong, Matthew Deans, Richard C Elphic, and David Wettergreen. Physical Process Models for Improved Rover Mapping. *iSAIRAS*, 2016.
- [10] Greydon T Foil. *Efficiently Sampling from Underlying Physical Models*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, October 2016.
- [11] Tom S. F. Haines and Tao Xiang. Active Learning using Dirichlet Processes for Rare Class Discovery and Classification. *IJCV 2013*, 2013.
- [12] Ju Won Hwangbo, Kaichang Di, Rongxing Li, and Geodetic Science. Integration of Orbital and Ground Image Networks for the Automation of Rover Localization. *ASPRS*, 2009.
- [13] D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *Journal Global Optimization*, page 13(4):455492, 1998.
- [14] a. Krause, C. Guestrin, a. Gupta, and J. Kleinberg. Near-optimal sensor placements: maximizing information while minimizing communication cost. *5th International Conference on Information Processing in Sensor Networks*, pages 2–10, 2006. doi: 10.1109/IPSNS.2006.244031.
- [15] David D Lewis and William A Gale. A sequential algorithm for training text classifiers. In *ACM SIGIR Conference on Research and Development in Information Retrieval*. Springer-Verlag New York, Inc., 1994.
- [16] Chun Kai Ling, Kian Hsiang Low, and Patrick Jaillet. Gaussian Process Planning with Lipschitz Continuous Reward Functions: Towards Unifying Bayesian Optimization, Active Learning, and Beyond. In *AAAI*, pages 1860–1866, 2016.
- [17] D. Lizotte. Practical Bayesian Optimization. *PhD thesis, University of Alberta, Edmonton, Alberta, Canada*, 2008.
- [18] Roman Marchant, Fabio Ramos, Scott Sanner, et al. Sequential Bayesian Optimisation for Spatial-Temporal Monitoring. In *UAI*, pages 553–562, 2014.
- [19] G. J. McLachlan and K. E. Basford. *Mixture Models: Inference and Applications to Clustering*. Marcel Dekker, 1988.
- [20] J. Mockus, V. Tiesis, and A. Zilinskas. *Toward Global Optimization*, volume 2. Elsevier, 1978.
- [21] L. Pedersen, M. Wagner, D. Apostolopoulos, and W.R. Whittaker. Autonomous robotic meteorite identification in Antarctica. *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation*, 4:4158–4165, 2001.
- [22] Carl Edward Rasmussen. *Gaussian Processes for Machine Learning*. Citeseer, 2006.
- [23] Nicholas Roy and Andrew McCallum. Toward optimal active learning through monte carlo estimation of error reduction. *ICML, Williamstown*, 2001.
- [24] William R Scott, Gerhard Roth, and Jean-François Rivest. View planning for automated three-dimensional object reconstruction and inspection. *ACM Computing Surveys (CSUR)*, 35(1):64–96, 2003.
- [25] P. Sebastiani and H. P. Wynn. Maximum Entropy Sampling and Optimal Bayesian Experimental Design. *Journal of the Royal Statistical Society*, 2000.
- [26] Burr Settles. Active learning literature survey. *University of Wisconsin, Madison*, 52(55-66):11, 2010.
- [27] Burr Settles, Mark Craven, and Soumya Ray. Multiple-instance active learning. In *Advances in neural information processing systems*, pages 1289–1296, 2008.
- [28] M. Shewry and H. Wynn. Maximum Entropy Sampling. *Journal of Applied Statistics*, 14(2), 1987.
- [29] Rr Sillito and Rb Fisher. Incremental one-class learning with bounded computational complexity. *Artificial Neural Networks ICANN 2007*, pages 58–67, 2007. ISSN 03029743.
- [30] D. R. Thompson, A. Allwood, C. Assad, D. Flannery, R. Hodyss, E. Knowles, and L. Wade. Adaptive sampling for rover x-ray lithochemistry. *i-SAIRAS*, 2014.
- [31] David R Thompson. *Intelligent Mapping for Autonomous Robotic Survey*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, August 2008.
- [32] David R Thompson, Greydon Foil, Michael Furlong, Anatha Ravi Kiran, and David Wettergreen. Spatio-spectral exploration combining in situ and remote measurements. *AAAI*, 2015.
- [33] et al. Thompson, David R. Automating X-Ray Fluorescence Analysis for Rapid Astrobiology Surveys. *Astrobiology*, 15:961–976, 2015.
- [34] T. J. Williams, B. J. and Santner and W. I. Notz. Sequential design of computer experiments to minimize integrated response functions. *Statistica Sinica*, 10: 11331152, 2000.