Robot Autonomy - Final Report

Spring 16

Information Theory Based Adaptive Exploration for Geological Classification

Team: Aishanou Osha Rait, Shivam Gautam, Tushar Agrawal, and Bishwamoy Sinha Roy

Mentors: Prof. David Wettergreen and Greydon Foil

Table of Contents

Problem Description

While exploring far-off planets it currently takes a lot of time to send precise waypoints to the rover due to inherent communication latencies. As such, current research methods focus on making rover exploration smarter. One approach is for the rover to pick waypoints in accordance to some optimization parameter.

The task at hand is the geological classification of a given region using noisy spectrometer data. First a satellite surveys a region of the planet and returns low resolution and noisy data of that region. Low resolution, in terms of spectroscopy, means the light broken down into a limited, often small, number of wavelengths. Rock classification based on this noisy data is unreliable. Hence, to supplement the satellite data a ground rover is provided with a high-resolution spectrometer which confirms or corrects the values of noisy data by physical sampling. However, using this spectrometer is resource costly, prohibiting continuous sampling. Consequently, the goal is to sample at selected locations so as to provide the cleanest data possible within the constraints of the resource budget.

Information gain is of importance in an information theoretic approach to adaptive exploration. The basic idea is to choose the key points so as to maximize the information gain from each sample [1]. The whole problem is inherently broken into two challenges. First, the rover needs to select key points, which is a global planner. Secondly, a local planner needs to plan a path to the key points, during which it can take multiple samples within a certain budget. As such, a plan for the rover not only contains waypoints but also sample points within the path.

In this project we focus on different ways of coming up with the waypoints and assume rover is sampling at each point. The algorithms are implemented on a simulated environment (details given in Appendix A) and not on the actual Zoe rover.

Challenge 1: Where to Sample

Approach 1: Differential Entropy

To maximize the information we want to sample at points where the entropy is the highest.
Differential Entropy is the method of determining entropy using the variance in the data values of different channels of spectroscopic data. It is given by [2]:

$$H(s) = \frac{1}{2} \sum_{b}^{B} \log(2\pi e \sigma_b)$$

Where σb is the variance in the bth band of S.
The higher the entropy more is the information gain after sampling that point.
The following steps outline the goal selection and updates to the path followed by the rover:

1. In the beginning the rover starts by sampling the starting location and calculates the entropy of all the points in the map w.r.t the starting point using the noisy satellite data. Thus the current sample set consists of only one point i.e., sample set S={S1}
2. Based on the entropy calculation, the highest entropy point is determined and the rover goes to that point and samples it. Thus the sample set increases, S={S1,S2} Entropy calculation will now include entropy between this sample set and all the remaining points on the map.

3. The rover continues to move towards the next highest entropy point while increasing its sample set, resulting in reduction of the maximum and average entropy as shown in figure 1 and 2 in Appendix B for 10 sets of randomly generated maps. This process stops once the sample budget is exhausted (taken as 50 samples).

Since this is a greedy approach, the robot follows a straight line path to each point, resulting in long and repetitive paths as shown in figure 1 below.

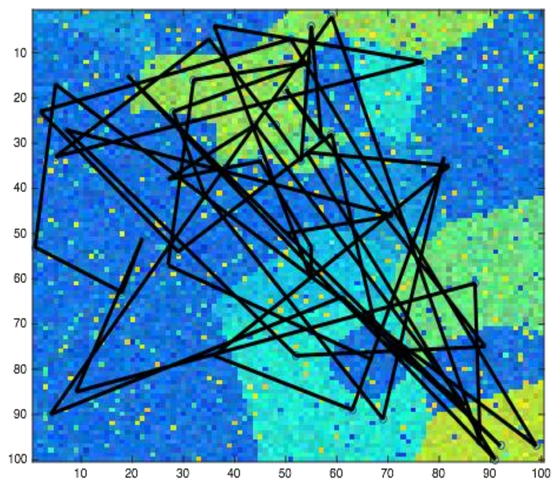

Figure 1 Path obtained through greedy sampling

To reduce this path length we tried to limit the search space by forming a window around current point thereby ensuring that the next sample point is within a certain distance from the rover's current location. Figure 3 in Appendix B shows that the paths obtained for a window size of +-20, which are much shorter and cleaner. However the entropy doesn't decrease uniformly over the entire test data (figures 4&5 of Appendix B). This is because the window is not able to capture the global information.

Through experimentation it was determined that a window of +- 40 results in reduced entropy with shorter paths. The comparisons for reduction in entropy per unit path length are shown in figures 7 and 8 of Appendix B. It can be concluded that there is no optimum solution for increasing information gain with reduced path length using the greedy approach. As a result more sophisticated techniques were evaluated as discussed further.

Approach 2: Shannon's Entropy via Clustering

The information theoretic approach to choosing a point to sample basically tries to optimize the amount of information gained by sampling, which is akin to decreasing "uncertainty" by the largest amount. To do this, we sample at the point that has the highest uncertainty. One way to formulate uncertainty was mentioned via differential entropy. Another way is using traditional Shannon's entropy, which is the expected information of a probability distribution. We defined the probability distribution over a set of discrete classes. Each class was defined by clustering the satellite image in feature space. Each wavelength in the satellite spectroscopic map was taken to be a feature. The data seems to be pulled from a mixture of Gaussian. Finding the cluster centers then is a Gaussian process.

We applied k-means and meanshift to find cluster centers. Since k-means is a parametric approach, we assumed the number of distinct classes of rocks in the map was known. On the other hand, meanshift doesn't assume number of classes but requires a variable called 'bandwidth' to be defined, which is basically how big the clusters should be in feature space, a variable that can be optimized. In a particular region, there exists two types of classes, dominant and rare classes. Since there are little points pulled from the rare classes, the classification via clustering will never find the rare class' cluster center.

Equation 1

$$P_{ij} = \frac{1/d_{ij}}{\sum_{j=1}^{J} 1/d_{ij}}$$

Equation 2

$$H(i) = \sum_{j=1}^{J} P_{ij} log_2(\frac{1}{P_{ij}})$$

Once the cluster centers are found, equation 1 is used to define a probability a point '**i**' falls within cluster '**j**' using the Euclidean distance, which ought to be modified to mahalnobis distance for appropriate use. Using this probability, Shannon's entropy is calculated using equation 2. Once the entropy is defined, the maximum entropy point is sampled. Sampling means pulling the spectrometer reading from the rover sample map (info 2 described in the simulation environment section). This moves the point in feature space because the true value of the point is now available. A Gaussian kernel is defined around the moved point to move nearby points. The standard deviation of the Gaussian kernel is proportional to the distance of the moved point to its cluster center. Basically, points that are very close to a cluster will create a larger Gaussian kernel. Without re-clustering, the equations are used to create an updated entropy map, from which a new sampling point is selected.

The algorithm was run on a test data of 10 different maps. The number of dominant and rare classes was varied, while keeping the total number of classes constant. Additionally, the probability of being picked from rare classes is varied as well. The results shown are very unstable. Although the entropy decreases by 1.25% for k-means and 2.96% in meanshift, as seen in the figure 2, the drop is very erratic.

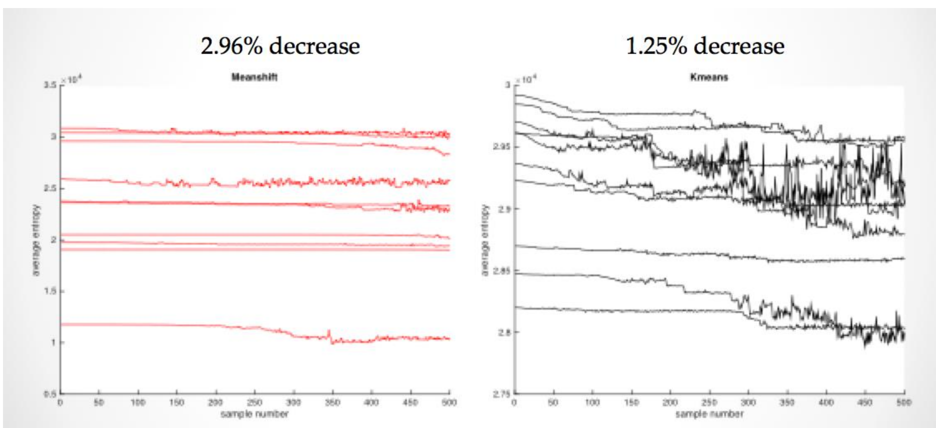

Figure 2: Average entropy from 10 test maps

Challenge 2: Planning a Path to sampling point

We discussed the problem of finding a goal for maximizing information till now. Now, given that goal, we want to find a path to the goal that optimizes information gain and path length.

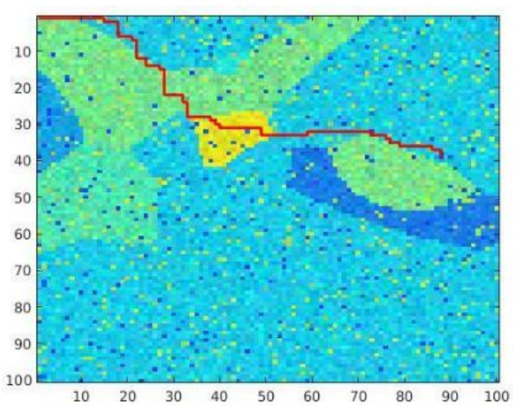

Figure 3 Path obtained from DP approach

Traditionally, scientists are assigned a budget of distance that can be used for exploration while travelling to the current goal. Then, scientists determine which rocks might be of enough interest while being inside the path budget. We find the best possible path for maximizing information given the constrained budget and use that as the gold standard to compare our heuristic approach.

Approach 1: Dynamic Programming

Dynamic programming is used in a breadth first search to find the path with maximum information gain, but for a constrained path length within the budget. This is then

used as a gold standard for comparing our heuristic approach. An example of a path is shown generated using the DP approach is shown in figure 3.
This path results in an information reduction of 69.75% with a 0.5449% entropy reduction per unit distance travelled. It should be noted however, that the runtime for this calculation was 5 hours, making it infeasible as a dynamic algorithm. The exact comparison can be seen in the statistical comparison section later.

Approach 2: Multi-Heuristic A*[3]

The Multi-Heuristic A* search algorithm addresses the issue that a single heuristic that captures all the specifics of a problem is hard to formulate. The solution proposed by this algorithm is to take into account an admissible heuristic in planning (the anchor heuristic) and multiple, arbitrarily admissible, heuristic functions.
For the purposes of our project, we decided to use the satellite information content as the additional inadmissible heuristic. This would allow us to combine the anchor heuristic of 'distance from goal' with the geological information of the map. This would, therefore, result in the expansion of nodes based on not only the anchor heuristic, but would also give priority to traversing through areas with high information.
The measurement units for the two heuristics are fundamentally different - the distance is in meters whereas the information of a particular cell is normalized to 0-1. To bring them to comparable terms, we used a factor, alpha, which scaled the information appropriately in the units of the anchor heuristic.

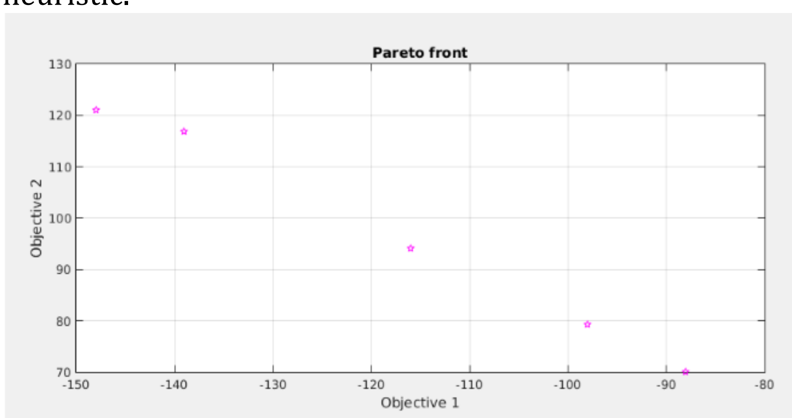

Figure 4 Pareto front

The value of alpha was set using a multi-objective Pareto optimization. The Objective 1 for the optimization was the distance travelled by the robot whereas Objective 2 captured the information content along the path. The values for the alpha parameter were noted and different paths were generated for each value. A qualitative comparison of paths generated based on information is described in the figures below. As can be seen from the results, if the information is given more weight, the robot tends to explore adjacent areas before reaching the goal. Conversely, if the robot does not take information into account, the path generated is similar to A* search.

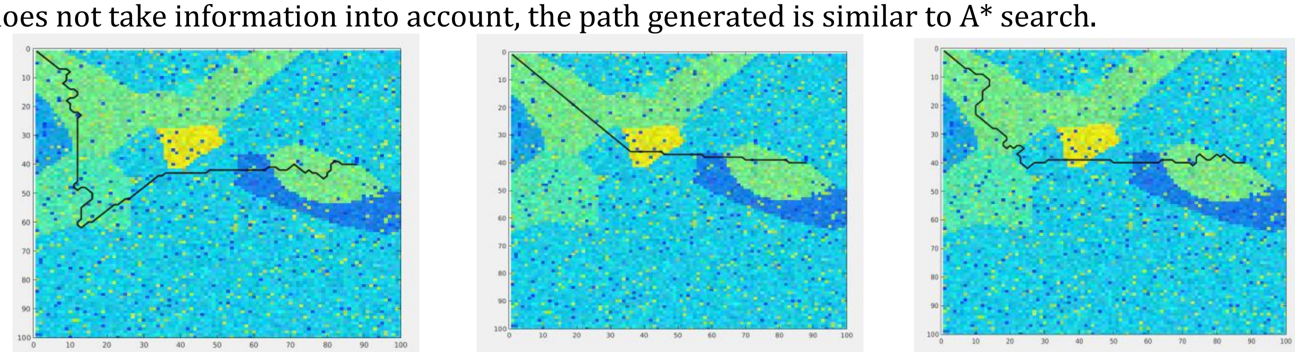

Figure 5 Paths generated based on different weights to information and distance. (left) Information>>Distance (center) Information<<Distance (right) Optimal weights

Statistical comparison of planning approaches

As is evident from above, the two approaches for planning are inherently better than the greedy best-first approach demonstrated in the case of searching in windows. A statistical comparison of the dynamic programming approach to the Multi-Heuristic A* approach is depicted in the table below. Both of them are also compared with the greedy approach for the full map and the windows.

S no.	Technique	% Entropy reduction	Entropy reduction per unit distance	% Entropy reduction per unit distance	Runtime
1	DP	69.75	211.66	0.5449	5 hours
2	MHA*	46.87	142.81	0.3982	0.1 s – 4 s
3	Greedy – Full map	100	3.5×10^{-4}	2.2×10^{3}	0.8 s -0.9 s
4	Greedy – Window Approach	100	5.6×10^{-4}	1.4×10^{3}	0.8 s -0.9 s

The multi-heuristic A* approach proved to be marginally inferior to the dynamic programming approach. The dynamic programming approach, even though statistically superior than the MHA*, took approximately 5 hours to find a solution for 100x100 grid as the number of states evaluated are exponential. Additionally, in the real world, the search would expand radially outward in the DP case which makes exhaustive enumeration over all states almost computationally intractable. The MHA* is able to come up with a comparable entropy reduction and takes only seconds to come up with a first solution.

Conclusion

The team was successful in emulating currently-used entropy sampling methods using the maximum entropy strategy. A novel approach using clustering based methods was explored and the groundwork for future research in this procedure was laid. In terms of planning, the team was successful in implementing planning using information with two different approaches. The approach which takes geological information content as a heuristic provided good results and this research provides a good base for future, possibly limited horizon, strategies in information-based planning.

References

[1]http://www.ri.cmu.edu/pub_files/pub4/thompson_david_r_2008_2/thompson_david_r_2008_2.pdf

[2] http://poincare.matf.bg.ac.rs/nastavno/viktor/Differential_Entropy.pdf

[3]Multi- Heuristic A* , Sandip Aine, Siddharth Swaminathan, Venkatraman Narayanan, Victor Hwan and Maxim Likhachev