

Ergodic Trajectory Optimization for Information Gathering

Kevin Edelson

October 31, 2020
CMU-RI-TR-20-55



The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

Thesis Committee:

David Wettergreen, *chair*
Maxim Likhachev,
Artur Dubrawski,
Alberto Candela

*Submitted in partial fulfillment of the requirements for the degree of
Master of Science in Robotics*

<https://youtu.be/a8OXHHGt2jA>

Abstract

Planetary robots currently rely on significant guidance from expert human operators. Science autonomy adds algorithms and methods for autonomous scientific exploration to improve efficiency of discovery and overcome limited communication bandwidth and delay bottlenecks. This research focuses on planning trajectories for information gathering and choosing sampling locations that have the most informative samples. We frame our exploration problem as a mapping problem for spectroscopic data and explore the concept of using low spatial density and low spectral resolution remote data as an information prior. We utilize a Gaussian Process regression model to fuse remote and *in situ* observations. This allows us to improve our high-resolution predictions across an entire scene without visiting all locations and compute an entropy map to guide exploration. We propose performing informative path planning using ergodic trajectory optimization. We explore the efficacy of the ergodic Spectral Multi-scale Coverage and ergodic Projection-based Trajectory Optimization algorithms. We demonstrate our approach in simulated exploration with real spectroscopic data of Cuprite, Nevada to highlight the advantages compared to traditional planning strategies. We successfully display that the ergodic Projection-based Trajectory Optimization planner outperforms all planners including state-of-the-art a non-myopic Markov Decision Process based planner. We also explore how the planning horizon affects planner performance, varying how often the entropy map is updated and the remaining sample points re-planned. We demonstrate that re-planning after every 1 sample does improve performance however a planning horizon of 4 samples offers a favorable balance between improved information gathering and increased computation time.

Acknowledgements

I want to express my gratitude to Professor David Wettergreen for his patient wisdom, his kind steady advice, and his love of the game (the game of robotics). I truly appreciate your commitment to advising your students as researchers as well as people. Thank you to Alberto Candela, who laid the foundation for this work and is a tremendous collaborator and even better person.

I want to thank my lab mates Srinivasan Vijayarangan, Himanshi Yadav, and Suhit Kodgule for their support and camaraderie. Thank you to the unsung heroes of the Robotics Institute, Suzanne Lyons Muth and Jean Harpley. Suzanne, you are a wizard, this department would come to a screeching halt without you, and you make it look easy. Jean, you consistently stick your neck out and go out of your way to help us students, RoboOrg owes you a debt of gratitude.

Endless love and appreciation to my dear friends Keene Chin, Mo Chin, Kate Shih, Aditya Dhawale, and Ada Taylor. Grad school can be a strange and intense experience and your support and having you as my chosen family will always make Pittsburgh feel like a home. I'm looking forward to the next chapters of our friendships (heart emoji).

Adam, I'm at Bed Bath Beyond and I can't remember, did you say to get metal or wood shower curtain rings? Joe, you either die a Buca or live long enough to see yourself become a Beppo. Gizzy, you've forgotten more than the entire human species will ever know. Rachel, a true mench, you inspire me to be a better scientist and a better musician. Ariel, I just wanted to say heyyyyy, and I'm here to stayyyyyy, and you gon' be mad everydayyyyyy, SUCCESS! Cookie and Chup, meow. Biddo, Clauds, and Craig, thank you for always having my back even when it's not easy for you, I love you.

This research was supported by the National Science Foundation National Robotics Initiative Grant #IIS-1526667, by the NASA Solar System Exploration Virtual Institute and the Planetary Sciences Institute project TREX: Toolbox for Research and Exploration (award 80ARC017M0005), and by the NASA Space Technology Fellowship (award 80NSSC19K1141).

Contents

1	Introduction	1
1.1	Science Autonomy	2
1.2	Problem Formulation	3
2	Spectroscopic Maps	5
2.1	Background	5
2.1.1	What is a Spectrum?	5
2.1.2	Hyperspectral Images	7
2.1.3	Spectral Features and Endmember Detection	8
2.2	Related Work	10
2.3	Active Spectroscopic Mapping	11
2.3.1	Spectroscopic Data Feature Extraction	12
2.3.2	Gaussian Processes for Spatio-Spectral Regression	13
3	Informative Path Planning	17
3.1	Related Work	18
3.2	Information Representation	20
3.3	Ergodic Trajectory Planning	21
3.3.1	Ergodicity Metric	21
3.3.2	Ergodic Spectral Multi-scale Coverage	23
3.3.3	Ergodic Projection-based Trajectory Optimization	26
4	Simulated Field Experiments	31
4.1	Data Sets	31
4.1.1	AVIRIS-NG	31
4.1.2	ASTER	32
4.2	Spectral Reconstruction Experiments	32
4.2.1	Experimental Design	33
4.2.2	Baseline Planners	34
4.2.3	Ergodic Planners	36
4.2.4	Planning Horizon - Model Predictive Control	38
5	Results and Analysis	40
5.1	Planner Performance Overview	40
5.2	Ergodic Planners	44
5.3	Variation Between Sites	46
5.4	Planning Horizon	52
5.5	Summary	54
6	Conclusion	56
6.1	Contributions	56
6.2	Future Work	57

A	Appendix	59
A.1	Site A	59
A.2	Site B	64
A.3	Site C	69
A.4	Site D	74
A.5	Site E	79

List of Figures

1	Prototype planetary rover in Cuprite, Nevada exploring terrain to classify and map geology. An on-board spectrometer measures ground spectra to identify mineralogical composition.	2
2	Example of a normalized reflectance spectrum.	6
3	Hyperspectral image of Cuprite, Nevada. Note the large number of channels per pixel compared to a standard RGB image. Two pixel location spectra are shown.	8
4	Spectral Linear Mixture Model: Exploded view of a pixel of an orbital image from a location in Cuprite, Nevada. The pixel consists of two endmembers, Alunite and Kaolinite, which linearly combine to form the observed spectrum.	10
5	The update step of the active spectroscopic mapping model. . . .	16
6	The prediction step of the active spectroscopic mapping model. .	16
7	Example Spectral Multi-scale Coverage trajectory on a simulated expected information density. Note that the amount of time spent in regions of the state space is proportional to the expected information density.	25
8	Example Projection-based Trajectory Optimization trajectory on a simulated expected information density. PTO is initialized with a trajectory computed from SMC shown in blue. Note the dynamics are still highly ergodic but the addition of a control effort cost allows us to better balance information gathering and path length.	30
9	The boundaries for each experimental site in the Cuprite, Nevada region. For site A we plan 10 samples, for sites B, C, and D we plan 20 samples, and for site E we plan 50 samples.	35
10	This example run shows typical path shapes that each planner achieves. Note the smooth paths of the ergodic planners that efficiently cover the region and the non-smooth MCTS path which backtracks. Random produces short, jagged, uninformed paths while Greedy often gets stuck in a local minimum somewhere on the map.	43
11	Average of all runs reported, lower entropy is better performance. Random performs the worst and performance improves as the planners become less myopic. PTO performs 31.2% better than Random and 16.3% better than Greedy in terms of entropy reduction. The two PTO initialization's performance is not statistically significantly different.	43
12	Ergodic planners produce smooth paths compared to baselines and therefore are good candidates for use with AUVs and UAVs.	45
13	Example SMC trajectory displaying multi-scale behavior. The path first visits locations with low-frequency information features and then priorities high-frequency features.	46

14	The statistical variation of final entropy between runs is shown here. There is a clear trend in median final entropy reducing as planning horizon reduces.	53
15	2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence.	53
16	Paths of multiple planning horizons from the same initial position. Notice the paths diverge from each other as the plans execute and as the planning horizon is shortened. The paths get less smooth and have more turns with the shorter horizons as the planners are seeing more localized bright entropy spots with more entropy map updates.	54
17	Greedy, as expected, is more susceptible to noise and gets stuck in a very sub-optimal minimum. All ergodic planners produce smoother paths compared to non-ergodic. PTO and MCTS globally look to cover similar points but the longer planning horizon of PTO allows it to perform best. Note the irregular shape of the MCTS path due to being fixed to the grid, MCTS also backtracks and repeats a sample location.	60
18	Examples like this show why Random often reports a high efficiency; even though the entropy score is always the worst, the paths are typically much shorter. MCTS again has a non-smooth path that crosses back over itself. Note that PTO-SMC and PTO-RAND move in different directions with different path shapes but consistently outperform baselines.	60
19	Average of all runs reported, lower entropy is better performance. Random performs the worst and performance improves as the planners become less myopic. SMC is competitive with MCTS despite SMC only planning single steps and MCTS has a look-ahead depth of 5 samples. Note that both PTO planners only perform best after nearly all samples are collected, this is expected as the planners optimize over the whole trajectory not each point individually.	61
20	Average of all runs reported, lower RMSE is better performance. Note that signal trends are highly correlated between entropy and RMSE, therefore our entropy calculation is a reliable information metric for use in informative path planning. Notably, the Greedy planner RMSE curve differs greatest from the entropy curve. It is the most myopic of the informative planners and susceptible to noise, i.e. it greatly favors exploitation over exploration. This suggests there that while entropy and RMSE are highly correlated there is noise on the entropy signal that must be accounted for.	61

21	The statistical variation of final entropy between runs is shown here. Greedy and SMC have similarly large IQR, this is expected given the single step planning of each. SMC is ultimately able to perform better than Greedy because the ergodicity metric it optimizes over is better able to balance exploration and exploitation. The planning horizon changes between planners does not appear to greatly effect statistical spread of planner performance.	62
22	2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence. MCTS and SMC variant performance are not different. Notably, most PTO variants are different than all other planners, while most PTO variants are not different than each other. PTO-RAND-INIT-MPC-1 performs best and is statistically different almost all other planners.	62
23	The statistical variation of final RMSE between runs is shown here. Once again, trends of the entropy performance are matched in the RMSE performance. The Greedy planner large IQR once again reflects the planner's susceptibility to noise.	63
24	2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final RMSE performance between planners with 95% confidence. The entropy performance again is reflected. The notable difference is that PTO-RAND-INIT-MPC-1 RMSE performance is not different than most other PTO variants.	63
25	Greedy again is stuck in a local minimum. The multi-scale nature of SMC is shown here, the path targets low frequency information features first before revisiting areas for higher frequency features.	65
26	Note how PTO is able to achieve nearly all the same points as SMC but without backtracking due to optimizing over the full trajectory. Given the current science goals moving on a to a new site may or may not be desired, however SMC is more likely to return to similar areas given its multi-scale nature.	65
27	Average of all runs reported, lower entropy is better performance. Random performs the worst and performance improves as the planners become less myopic. We once again see the PTO planners over taking others as the samples approach the full trajectory due to their being highly non-myopic.	66
28	Average of all runs reported, lower RMSE is better performance. Note that signal trends are highly correlated between entropy and RMSE, therefore our entropy calculation is a reliable information metric for use in informative path planning.	66

29	The statistical variation of final entropy between runs is shown here. Greedy has a particularly large IQR indicating there are local minima that myopic planners are susceptible to. SMC's multi-scale nature allows it to avoid the noisy local minima. There is a clear trend in median final entropy reducing as planning horizon reduces.	67
30	2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence. MCTS is not statistically different than a few ergodic variants but is notably different than the best performing planner, PTO-SMC-INIT-MPC-1. The SMC variants are not statistically different from each other as are many of the PTO-SMC variants.	67
31	The statistical variation of final RMSE between runs is shown here. Once again, trends of the entropy performance are matched in the RMSE performance, particularly when comparing planning horizon within a planner type. Greedy notably has a much smaller IQR for RMSE compared to entropy.	68
32	2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final RMSE performance between planners with 95% confidence. The entropy results are reflected here with notable differences. The best performing planner, PTO-SMC-INIT-MPC-1, is no longer statistically different than its competition, namely MCTS. This is more a function of the spectroscopic mapping model as the planners directly optimize over the entropy signal.	68
33	The smoothness of ergodic planner paths is shown well here. The MCTS path by comparison has many sharp turns and crosses back over itself. This example shows the Greedy planner performing well but the sharp back and forth about half way into the path shows the sensitivity of the planner to noise.	70
34	This example highlights the difference between PTO-SMC-INIT and PTO-RAND-INIT. Each planner converges to a different minima despite the same initial position however both are very effective at reducing entropy.	70
35	Average of all runs reported, lower entropy is better performance. Random performs the worst and performance improves as the planners become less myopic. Site C shows SMC not performing as closely to the best planners compared to the other sites. . . .	71
36	Average of all runs reported, lower RMSE is better performance. Note that signal trends are highly correlated between entropy and RMSE, therefore our entropy calculation is a reliable information metric for use in informative path planning.	71
37	The statistical variation of final entropy between runs is shown here. There is a clear trend in median final entropy reducing as planning horizon reduces.	72

38	2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence. PTO-SMC-INIT-MPC-1 performs best and is not statistically different than high performing planners MCTS, PTO-SMC-INIT-MPC-1, and PTO-SMC-INIT-MPC-4.	72
39	The statistical variation of final RMSE between runs is shown here. Trends of the entropy performance are matched in the RMSE performance.	73
40	2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final RMSE performance between planners with 95% confidence. The entropy performance again is reflected. The statistical differences are similar only variants that differ by planning horizon are no longer statistically different.	73
41	Site D is highly uniform in spectral diversity. The large white location is homogeneous in mineral type and therefore spectra signature. Note that planners effectively navigate the bi-model distribution. This site was chosen to investigate performance on regions with more uniform belief state, i.e. a less informative prior.	75
42	This example shows the planners spending more time in the grey locations with more information. PTO and SMC appear to more favorably balance exploration and exploitation but not aggressively avoiding the white location like MCTS.	75
43	Average of all runs reported, lower entropy is better performance. The performance difference between all planners is more pronounced in site D due to the more uniform information distribution. Note that all ergodic planners outperform other baselines indicating they do favorably balance exploration and exploitation.	76
44	Average of all runs reported, lower RMSE is better performance. Note that signal trends are highly correlated between entropy and RMSE, therefore our entropy calculation is a reliable information metric for use in informative path planning.	76
45	The statistical variation of final entropy between runs is shown here. Greedy has a large IQR given the myopic nature of the planner and the uniformity of the site. Note that the PTO-SMC-INIT variants contain more outliers than usual due to local minima in the uniform site.	77
46	2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence. Nearly all of the PTO variants are statistically different from other baselines and outperform them including MCTS.	77
47	The statistical variation of final RMSE between runs is shown here. Once again, trends of the entropy performance are matched in the RMSE performance.	78

48	2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final RMSE performance between planners with 95% confidence. The entropy and RMSE statistical difference is nearly identical with only PTO-SMC-INIT-MPC-8 no longer different than SMC-MPC-1 and PTO-SMC-INIT.	78
49	Only the baseline planner paths are shown here due to the long paths. Greedy is distracted by noise as is expected.	80
50	The ergodic planner paths are shown here. Their paths are notably smoother than the baseline planner paths. There is path crossover in the PTO-SMC-INIT path but the curves remain smooth. The SMC has sharp turns which is explained by significant changes to the entropy map after the previous sample.	80
51	Average of all runs reported, lower entropy is better performance. Random performs the worst and performance improves as the planners become less myopic. Of significance here is SMC's strong performance, outperforming MCTS throughout the entire average of the trajectories.	81
52	Average of all runs reported, lower RMSE is better performance. Note that signal trends are highly correlated between entropy and RMSE, therefore our entropy calculation is a reliable information metric for use in informative path planning.	81
53	The statistical variation of final entropy between runs is shown here. Interestingly PTO-SMC-INIT-MPC-8, PTO-RAND-INIT-MPC-8, and SMC-MPC-8 have large IQR or outliers while the other planning horizon variants do not. This suggests there are local minima modes in the entropy mode that are highlighted by updates after 8 steps but not finer or courser updates.	82
54	2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence. The majority of planner variants are statistically different from one another. Notably, the best performing planner PTO-SMC-INIT-MPC-1 is statistically different from all others except SMC-MPC-4 and PTO-RAND-INIT-MPC-1.	82
55	The statistical variation of final RMSE between runs is shown here. Once again, trends of the entropy performance are mirrored in the RMSE performance.	83
56	2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final RMSE performance between planners with 95% confidence. Key differences between entropy performance and RMSE performance include the high performance of SMC variants. They are no longer statistically different from the highest performing planner, PTO-SMC-INIT-MPC-1.	83

List of Tables

1	Spectrometer Instrument Comparison	8
2	Planner Property Comparison	38
3	Best Performing Planner by Site	40
4	Site A Planner Comparison	47
5	Site B Planner Comparison	48
6	Site C Planner Comparison	49
7	Site D Planner Comparison	50
8	Site E Planner Comparison	51

1 Introduction

Modern planetary robotic exploration is guided by scientists specifying the locations on a path that they believe will best address mission investigation questions. The path is formed from expert knowledge of the site and expectations about where to gather scientific information. Scientists reinterpret their measurements with growing contextual knowledge of the environment, so exploration is characterized by a frequent reformulation and replanning throughout the mission lifetime. Replanning occurs on large strategic scales, bypassing or favoring geographic locales, as well as local tactical scales, to approach potential new discoveries or to remain at an anomalous feature for additional measurements. However, many exploration scenarios occur with low bandwidth and high latency communication, leaving limited opportunities to revise exploration plans.

This work describes an approach to overcome the communication bandwidth and delay bottleneck in robotic exploration, where the command to the remote explorer is based on an evolving model of the explored environment, rather than a single prescribed route, enabling the robot to take more adaptive and efficient actions based on real-time information, improving the rate and productivity of discovery. By encoding belief of the spatial distribution of the information content of a region, we can iteratively plan actions that will guide the robot toward more information rich areas. Our context driven exploration strategy draws inspiration from and captures the characteristics of geologic site survey. Sample interpretations are sensitive to spatial relations. Rather than considering each sample independently, field scientists incorporate each new observation into an evolving model of geologic formation.



Figure 1: Prototype planetary rover in Cuprite, Nevada exploring terrain to classify and map geology. An on-board spectrometer measures ground spectra to identify mineralogical composition.

1.1 Science Autonomy

Algorithms and methods for autonomous exploration and scientific measurement by robotic agents falls under the label of Science Autonomy. In this work, we are focusing on information gathering in the context of spectroscopic mapping, which enables analysis and interpretation of material composition and physical properties. Our objective is to discover the true state of some phenomena of interest from the true spectral reflectance of a region. Therefore, the robotic explorer should collect observations that are most informative with respect to the spatial model, the spectroscopic map. Exploration becomes an active learning problem, where a data sample's value is derived from the information content it provides to the larger spectroscopic map.

Often regions of interest are studied via airborne or orbital remote instruments prior to deploying an *in situ* robotic instrument. Remote sensing in-

struments often suffer from low spatial density and low spectral resolution and therefore robots are tasked with collection of high resolution data. We limit our exploration to the common scenario where low resolution data is available for a region prior to the deployment of the *in situ* robotic explorer. Ultimately the goal is to integrate both types of measurements to reconstruct dense maps of high resolution spectra without taking high resolution measurements of all map locations. Therefore it is imperative to plan robot trajectories that contain the most information to efficiently reconstruct spectroscopic maps. We demonstrate the feasibility and successful performance of our high-level autonomy approach via simulation investigations of real data from Cuprite, Nevada.

Our evaluation of the efficacy of our robot trajectories are based on comparison of reconstructed spectra with ground-truth high resolution real data measurements. This performance metric, based on mapping accuracy, inherently ignores some aspects of exploration problems. Novel feature detection is disregarded, as the underlying phenomena to be mapped is predefined. Also, our formulation does not address the needle in a haystack problem.

1.2 Problem Formulation

Our specific formulation of this active exploration problem includes leveraging prior data about a region of interest to guide autonomous robot exploration to ultimately produce a high resolution spectroscopic map. Low spatial density and low spectral resolution data provides significant information to guide our robot exploration compared to no prior information (uniform prior). We encode the spectroscopic data into a latent space Gaussian Process representation. Thus we are able to reduce the dimensionality of *in situ* samples to the same number of features as our prior data to integrate both information sources.

Our formulation of belief space representation, a Gaussian Process, is a con-

tinuous model which provides the advantage of not requiring discretization (or arbitrary discretization) of the space to plan trajectories. Our contributions therefore include the study of ergodic trajectory optimization to plan informative robot trajectories. Ergodic trajectory optimization allows for a continuous information space and continuous arbitrary non-linear agent dynamics which provide distinct advantages over previous methods. Often the map, belief space, and agent dynamics are required to be discretized to employ search-based and Markov Decision Process based trajectory planning. Continuous space planning provides the inherent benefit of allowing the robot to take samples not represented in a grid, therefore favorably balancing exploration versus exploitation. We will compare the performance of ergodic planning techniques with baseline informative search techniques in terms of computation and the sample efficiency of each trajectory. We will explore the usefulness and robustness of ergodic Projection Based Trajectory Optimization (PTO) for users to balance the relative cost of maximizing information gain and minimizing control effort. We also will explore the use performance of ergodic planners in a Model Predictive Control (MPC) framework to better understand how planning horizon and re-planning affects trajectory sample efficiency. We therefore contribute methods for applying various ergodic trajectory optimization techniques for a spectroscopic mapping application, implement the algorithms from scratch, perform simulated experiments on real data to interrogate the performance of the ergodic techniques, and analyze results to make claims about the best uses of ergodic trajectory optimization.

2 Spectroscopic Maps

In short, spectral signals allow us to infer what something is made of. The spectral signature of a surface provides insight into its composition and subsequently its formation process. Planetary scientists use imaging spectrometers to measure the light reflected from a surface at various wavelengths. The natural process of material absorbing and reflecting light differently at different wavelengths allows scientists to infer the composition of the target with a spectral signal. The number of wavelengths an instrument measures with each sample is referred to as the number of channels. More channels (i.e. more wavelengths) allow greater insight into the spectral features of a spectra. Spectral features are the signal shapes that allow scientists to infer composition from a signal.

Imaging spectrometer instruments are commonly deployed on spacecraft and aircraft to map the spectral content of regions. Various instruments are described in detail in Section 4.1.1 and Section 4.1.2. These remote measurements lack in spatial and spectral resolution however, and are often insufficient to planetary scientists to make confident predictions of surface composition and material distribution. Our method allows for *in situ* robotic explorers to leverage this prior low-resolution information to guide exploration. The *in situ* and remote measurements can also be fused to better predict high-resolution spectral signals throughout a scene.

2.1 Background

2.1.1 What is a Spectrum?

We refer to a series of reflectance values as a spectrum. Reflectance is a measure of the electromagnetic energy present in the light reflected or scattered from a surface at a particular wavelength. More photons present, yields a higher reflectance measurement. A spectrum refers to a series of reflectance measure-

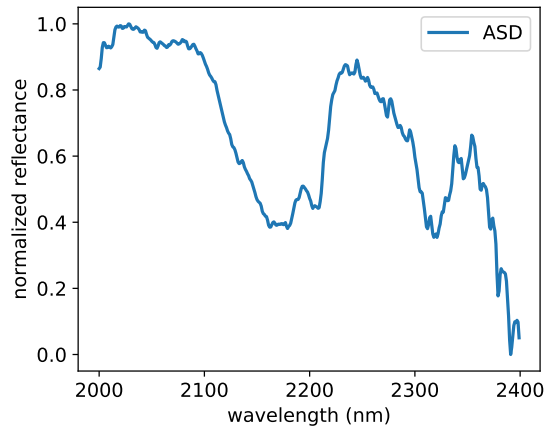


Figure 2: Example of a normalized reflectance spectrum.

ments from the same source but at different wavelengths of light. A spectrum captures a significant amount of information about the source. The pattern of reflection and absorption bands allows scientists to derive information about the chemistry and composition of a material from the light reflected from the material. As photons enter a material, some are reflected from surfaces, some pass through the surface, and some are absorbed. Those photons that are reflected from surfaces or refracted through a particle are said to be scattered.

Spectra are measured using a variety of instruments but we will briefly discuss how one particular class of instruments works, the spectrometer. This allows us to better understand the physical process of the taking measurements and how that effects the formulation of our trajectory planning problem. As you might assume, the basic function of a spectrometer is to measure a spectrum. That is, to take in reflected light from a target, break it into its spectral components, and digitize the signal as a function of wavelength. Measurements are normalized to compensate for changes in the incident light source, in our case changes in sunlight conditions. Therefore, spectrometer measurements are taken of a white reference material, which provide spectra with which to normalize

subsequent measurements with the same lighting conditions.

2.1.2 Hyperspectral Images

Our spectroscopic maps leverage the data structure of a hyperspectral image. A hyperspectral image is like a standard RGB image, but instead of each pixel containing an intensity value for three color channels there are more channels that refer to specific wavelength bands and can also include wavelengths outside of the visible spectrum. Instruments vary in terms of the number of wavelength channels, the width of each wavelength band, and what part of the electromagnetic spectrum is covered. Different wavelengths are more informative for detecting different materials and often the wavelength range used is chosen based on estimates of material composition. The number of channels in each spectra is referred to as the spectral resolution. Consequently, more channels provide better insight into the true spectral signal and is therefore more informative in identifying material composition. The spatial resolution of the hyperspectral image, the physical distance measured by each pixel, is a function of instrument optics and relative instrument and target positions. Imaging spectrometer instruments are often used in orbital, airborne, and ground-based settings. A hyperspectral image from an instrument used on orbit will, on average, have a lower spatial resolution than one from a ground-based instrument. The geometry of the position of the orbital instrument and the optics field of view will yield pixels that cover much larger physical regions than ground-based instrument. It follows that there is a primary issue when using low spatial resolution data to determine material composition. How can one identify the composition of a target if its spectra contains a diversity of materials?

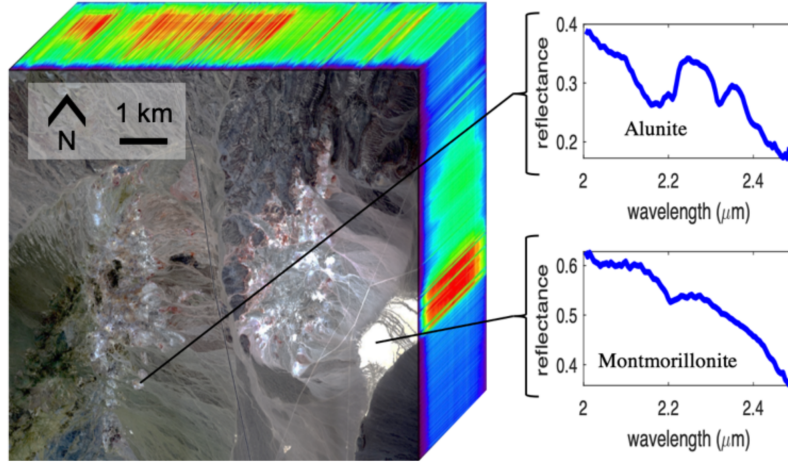


Figure 3: Hyperspectral image of Cuprite, Nevada. Note the large number of channels per pixel compared to a standard RGB image. Two pixel location spectra are shown.

Table 1: Spectrometer Instrument Comparison

Instrument	Source	Wavelength Range	Channels	Avg. Res.
ASTER	Orbital	2-2.4 μm	5	80 <i>nm</i>
AVIRIS-NG	Airborne	2-2.4 μm	80	5 <i>nm</i>
ASD	Ground	2-2.4 μm	400	1 <i>nm</i>

2.1.3 Spectral Features and Endmember Detection

Measuring the spectral signature of materials is of interest to scientists to determine material composition. A reflectance spectrum contains unique signal features that allow scientists to very accurately identify the sampled material. When a stream of photons encounter a medium with a change in the index of refraction some are reflected and some are refracted into the medium [13]. This results in absorption patterns, where the relative strength of reflectance and refractance across various wavelengths creates unique signal features dependent on material [46].

Spectroscopists often classify material composition based on spectral measurements using reference measurements of pure laboratory materials, this is called endmember detection. The Tetracorder expert system, developed by Clark et al., uses an extensive library of USGS spectra to identify endmembers [12]. Other times learning-based approaches at classification of spectra are used [62]. In all cases a set of reference measurements are required and your classification is only as good as your reference measurement signals. Therefore we focus on reconstructing high-resolution spectra and decouple the problem of material classification from our problem of information gathering.

Spectral signals often are combinations of endmembers, a phenomena known as spectral mixing. If the incident light into the spectrometer instrument reflected from various materials, the resultant signal will contain spectral features of all materials present proportional to the concentration of each material. Figure 4 show the concept of spectral mixing. This adds another layer of complexity to classification directly from our GP model. There is a mature knowledge base regarding spectral mixture models and methods for unmixing spectra for the purpose of endmember detection [43]. We therefore choose to focus on reconstructing high-resolution spectra and using other means for unmixing and endmember detection.

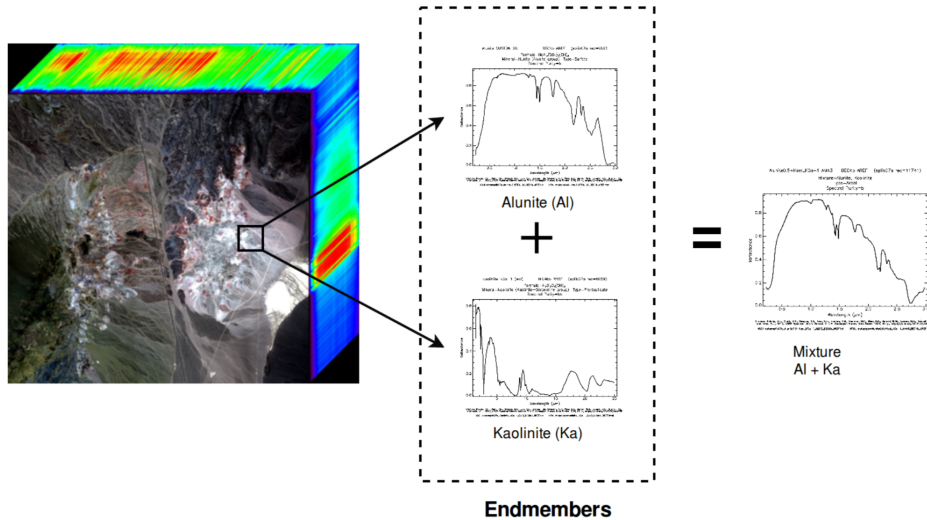


Figure 4: Spectral Linear Mixture Model: Exploded view of a pixel of an orbital image from a location in Cuprite, Nevada. The pixel consists of two endmembers, Alunite and Kaolinite, which linearly combine to form the observed spectrum.

2.2 Related Work

Related work provides models for representing spectroscopic maps. Thompson’s work [57] specifically developed spatial models of the environment: a generative map of the environment that extrapolates from previous measurements and predicts new measurements at un-visited locations. The work demonstrated an effective spatial model representation using Gaussian Processes that infused orbital data as input dimensions into the model. This Gaussian Process spatial models allow for adaptive, online learning such that an agent can efficiently update the models in real-time. A distinct difference of this work from robotics concepts like topological mapping is that unlike spatial modeling, topological mapping does not extrapolate beyond the sensing-horizon of the robot [20]. One of the reasons why spatial modeling is able to achieve this is because of strong cross-sensor correlation such as between orbital data and *in situ* data.

For example, geological studies of the reflectance spectra measured by these spectrometer instruments can characterize the minerals present on the surface. Therefore, comparing the spectral signatures of *in situ* measurement with the corresponding orbital spectra at a particular location allows us to make informed predictions about the minerals present at locations where similar orbital spectra are observed.

Adapting concepts from information-driven sampling for sensor networks [33], Thompson further proposed information gain based sampling techniques for spatial modeling. Building upon Thompson’s work, Foil [17] used an active learning approach to expand on sequential modeling. Sequential modeling is a topic similar to spatial modeling, with a key difference being that sequential modeling allows non-stationary environmental models. Foil also proposed a Dirichlet Process based Adaptive Gaussian Mixture Model model and a Gaussian Process based model for adaptive sampling; a form of sequential modeling with non-stationary objective functions. Recently, Candela et al. [7] developed the science hypothesis map: a probabilistic structure in which the initial beliefs of scientists evolves using Bayesian updates as the robot takes measurements.

2.3 Active Spectroscopic Mapping

Our goal is to leverage data collected from low spectral resolution and low spatial resolution sources in order to plan trajectories that collect high spectral resolution and high spatial resolution measurements to best reconstruct a spectroscopic map. The question remains: how can we mitigate the resolution disparity and combine these sources of data in a spatial model? We also would like to be able to iteratively update the our spatial model with each new measurement taken. It stands that our model should permit online learning and inference while representing both spatial and spectral correlations in collected

data. We rely on Gaussian process regression, a machine learning technique that has been widely used in spatial statistics [15] and informative robotic exploration [5]. We opt to use a Gaussian Process Regression based spatial model derived from Thomson [57] combined with feature extraction methods from Candela [8].

2.3.1 Spectroscopic Data Feature Extraction

Gaussian Processes (GPs) are typically used for representing scalar-field values [47], however spectra are inherently high dimensional. *in situ* spectroscopic measurements tend to have very high spectral resolution, which makes using a GP for each measurement channel computationally infeasible. However, many of these channels and wavelength bands are highly correlated, allowing for the application of dimensionality reduction techniques. We use a Variational Autoencoder (VAE) [30] technique developed for spectral application by Candela [9]. A VAE is a neural network that converts a set of high-dimensional observations $y \in Y \subset \mathbb{R}^n$ into a set of lower dimensional features $z \in Z \subset \mathbb{R}^d$, where $d < n$. We specifically use a VAE because it learns a representation that resembles a multivariate normal distribution, i.e. $Z \sim \mathcal{N}_d(0, I_d)$, effectively normalizing and uncorrelating the features. The VAE uses the following loss function:

$$L(X) = Error(\hat{X}, X) + \lambda D_{KL}(\mathcal{N}(\mu, \Sigma), \mathcal{N}(0, 1)) \quad (1)$$

with the first term referring to the reconstruction error and the second term is a regularizer that refers to the amount of information lost in the compressed representation of the data [30]. The parameter λ is a tunable gain to weight the relative importance of the reconstruction accuracy and the size of the compressed representation. The VAE is comprised by two networks: an encoder that extracts the features, and a decoder that reconstructs high-resolution ob-

servations using the learned features. It is important to underscore that the VAE ignores spatial correlations in the data. It is also important to note that the low dimensional features output by the VAE do not represent spectral features in the sense of the spectroscopy. In other words, there is no physical interpretation of the features, but they allow for computationally efficient use of GPs to represent spectroscopic maps.

2.3.2 Gaussian Processes for Spatio-Spectral Regression

We use Gaussian Process regression to learn the spatial distribution of spectra in a spectroscopic map, \mathcal{M} . This process is defined as spatio-spectral regression. We follow from Candela’s [8] work, combining different sources of spectral data of varying resolutions. To mitigate the issue of high dimensional spectral data, we use GP regression to learn distributions in the condensed feature space, $Z \subset \mathbb{R}^d$, from the VAE. The VAE dimensionality reduction allows data sources of disparate dimensionalities to be reduced to a common size and furthermore uncorrelates the learned feature representation, allowing for the use of d independent GPs.

We assume we have low resolution remote spectra $x \in X \subset \mathbb{R}^m$, and high resolution *in situ* spectra $y \in Y \subset \mathbb{R}^n$. We assume remote spectra are available prior to deployment of the robotic explorer for the spatial locations of interest $l \in L \subset \mathbb{R}^2$, whereas *in situ* spectra are collected by the robot in the field. We define the GP independent input vector variable v to be the known properties of the measurement location. For example, v could include the site’s physical position, represented by latitude and longitude coordinates. It could also include other input dimensions corresponding to preexisting measurements or remote sensing data. We use the latitude and longitude position of the center of our remote spectra $l = [lat, lon]$, along with m remote spectral reflectance values $x = [c_1, c_2, \dots, c_m]$ and append them to produce the input vector $v = [l, x] \in$

$V \subset \mathbb{R}^{2+m}$.

We assume there exists a function $f^i : \mathbb{R}^{2+m} \rightarrow \mathbb{R}^1$ that maps an input vector v to each feature of vector $z \in Z \subset \mathbb{R}^d$, therefore $z_i = f^i(v) + \epsilon_i$ for $i = 1, 2, \dots, d$. Each of the d GPs learns a distribution over the the values of function f^i . Each GP is defined by a mean function μ_i and a covariance function K_θ^i [47]. We assume that the mean is zero, which is reasonable because of the nature of VAE normalized features. For the covariance matrix, we utilize a radial basis function (RBF) kernel defined as follows:

$$K_\theta^i(v, v') = \theta_0^i \exp \left(-\frac{\|l - l'\|_2^2}{2(\theta_l^i)^2} - \frac{\|x - x'\|_2^2}{2(\theta_x^i)^2} \right) \quad (2)$$

where $\theta^i = (\theta_0^i, \theta_l^i, \theta_x^i)$ are the kernel hyperparameters for each GP. We also use the GP variant for noisy observation and therefore add the additional noise hyperparameter, σ_{noise}^i [47]. The GP hyperparameters are estimated by maximizing the log-likelihood of the observed data as shown in [47]. We currently pre-tune the GP hyperparameters using a portion of the relevant dataset and keep these hyperparameters fixed during experiments, however the hyperparameter tuning process can be done iteratively with each new measurement.

The full active spectroscopic mapping model integrates VAE feature extraction and GP regression. Given a set of *in situ* observations, we would like to predict the true value of $f^i(v)$ at other locations not visited by the robotic explorer. We can infer a probability distribution from a prior $P(f^i(v))$ using Bayes' rule and Recursive Bayesian estimation. The probability distribution over possible $f^i(v)$ therefore provides a map of any desired resolution. This is done by evaluating Maximum A Posteriori (MAP) estimates of $f^i(v)$ at unobserved remote sample sites to yield the most likely observations [57]. The Recursive Bayesian estimation learning process is comprised of an update step and prediction step [59]. In the update step, shown in Figure 5, the learning model is improved

when the robot collects *in situ* measurements. The d independent GPs are updated using the compressed features that are extracted using the encoder of the VAE, along with the associated spatial coordinates and remote measurements. In the prediction step, shown in Figure 6, the model uses this new information and its correlations to old information to better reconstruct spectra in all map locations. First, the GPs predict the compressed features at each point of the map, $v \in V$, using a normal distribution, $\hat{Z} \sim \mathcal{N}(\hat{\mu}_Z(V), \hat{\Sigma}_Z(V))$. These predicted compressed features are then passed through the decoder of the VAE and reconstructed as high resolution spectra, \hat{Y} , that can be used to make more accurate predictions of material composition.

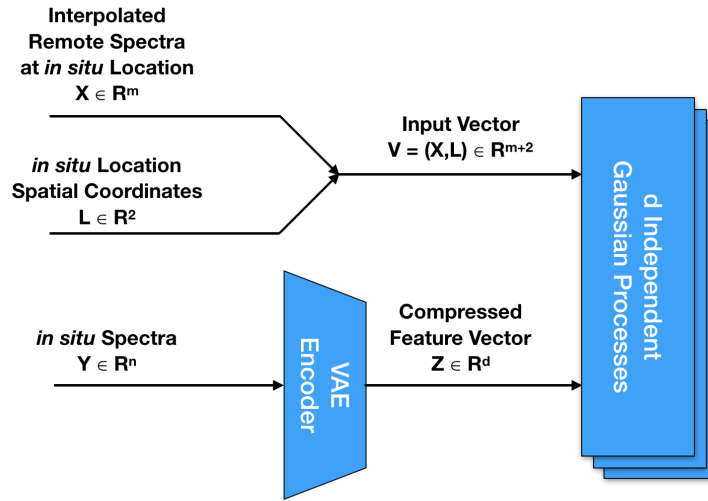


Figure 5: The update step of the active spectroscopic mapping model.

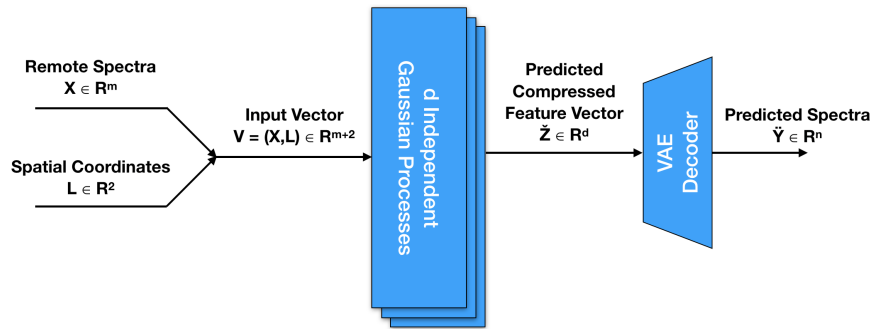


Figure 6: The prediction step of the active spectroscopic mapping model.

3 Informative Path Planning

Planetary robotics research has historically been focused on risk-aware planning; generating safe trajectories with an emphasis on hazard detection and avoidance and formulating risk-bounded temporal plans. However, with the advancement of space hardware capabilities [14] and the expectation for planetary rovers to explore terrains with a greater degree of autonomy [61], research on integrating higher-level science goals into the robot planning has risen in significance. The field of Science Autonomy focuses on developing robotic technologies that improve a planetary rover’s interpretation of the environment and subsequently makes decisions to maximize certain science objectives.

While work has been done in improving environmental models for science autonomy, research on improving the planning algorithms that would directly use these models for determining optimal sampling locations is still in its early stages. Planners currently either rely on science-agnostic strategies or myopic strategies that only attempt to maximize the immediate information content. We now examine the field of Informative Path Planning and discuss how techniques developed there can be translated to science autonomy.

The problem of guiding a robot autonomously, while intelligently selecting sample locations is referred to as Informative Path Planning (IPP). This problem may seem very similar to sequential modelling and indeed, some of the applications of sequential modelling can be brought under the umbrella of IPP problems. However, traditionally IPP has only focused on scalar sensory measurements like the temperature, salinity, or chlorophyll content of a body of water [5] or the radio frequency signal intensity of a region [11]. We introduce the additional nomenclature of informative path planning in order to distinguish it from sequential modelling that also considers high-dimensional inputs such as the spectra. Note that we will be forgoing a background on the basics of robot

path planning, given the which can be readily found in [29].

3.1 Related Work

The concept of Informative Path Planning (IPP) was introduced by Binney et al. [6] who used an exhaustive search strategy that exploited the monotonicity of certain objective functions for faster convergence. Lim et al. [35] developed Recursive Adaptive Identification (RAId), a polynomial-time approximation adaptive IPP algorithm that chooses the next sample location based on all information acquired at time of computation. RAId was tested on a number of tasks, such as robot grasping a 2-star graph search, and was shown to outperform baseline policies in these tasks. Cao et al. [10] proposed two multi-agent IPP algorithms based on entropy and mutual information as objective functions and demonstrate that they perform better than previous state-of-the-art algorithms with increasing planning horizons. Gautam et al. [21] employs an exploration strategy based on Multi-Heuristic A^* , to handle the trade-off between exploration cost and information gain through Pareto-optimal solutions. Morere et al. [41] developed a Bayesian Optimization based Partially Observable Markov Decision Process framework used for finding approximate solutions to trajectories to map scalar fields. Kodgule et al. [32] extended the use of Markov Decision Processes for IPP using a Monte Carlo Tree Search method to plan trajectories for spectroscopic mapping. Each of these IPP techniques relies on prior discretization of either the robot state space, the robot action space, the information state space, or all of the above. There are inherent losses in the exploration vs. exploitation trade-off when discretizing information state spaces, and the discretization of robot state and action space adds complexity depending on robot dynamics.

IPP algorithms that are able to operate in continuous spaces provide a richer

set of possible trajectories for monitoring the environment. Hollinger et al. [27] extended sampling-based methods, combining ideas from asymptotically optimal rapidly exploring random trees (RRT*), rapidly exploring random graphs (RRG), and probabilistic roadmaps (PRM*) with insights from branch and bound optimization. These strategies have been used for information gathering and mapping of scalar fields with success but requires the re-sampling of the information space with each new sample if an adaptive IPP approach is to be used.

Optimization-based informative planning methods integrate the fields of ergodic theory and optimal control. The use of ergodic coverage concepts for IPP problems stems from the inherent submodularity of information [44]. Simply, if you have covered and measured all states then you have all available information. Ergodic trajectory optimization has not been exhaustively researched in information gathering contexts and so far has only been used to localize the positions of unknown objects [39], [51], or map a scalar field [1], [3], . Optimization-based planning also benefits from operating in continuous spaces and allows for tight integration with robot dynamics models. This permits the use of arbitrary non-linear robot dynamics in many planner variants [38], [16], [3] and planners are capable of directly computing control trajectories rather than waypoints or motion primitives like search-based or sampling-based planning. Ergodic optimization-based planners often extend naturally into multi-agent tasks [4], [60], due to the formulation of the ergodicity metric that will be expanded on below. We will be focusing on single agent planning in this work to keep the scope in line with planetary robotics. We expand on the use of ergodic trajectory optimization to plan trajectories for spectroscopic mapping and compare against baseline methods for IPP.

3.2 Information Representation

All Informative Path Planning (IPP) techniques require the quantification of information content in order to incorporate an informative metric into an objective function. Integrating entropy or mutual information metrics in objective functions has been used in several IPP works with success [21], [11], [36]. Entropy is a measure of the uncertainty, therefore is directly correlated with expected information. In a spectroscopic mapping context, Kodgule [32] and Candela [8] demonstrated a strong correlation between entropy reduction and decreasing spectral reconstruction error. Consequently, we utilize entropy based objective functions in our approach to ergodic trajectory optimization.

We leverage the Gaussian Process (GP) representation of our spatial model to efficiently compute spatial entropy and produce an entropy map. Our use of a variational autoencoder (VAE) to reduce all spectral data into an uncorrelated feature space of dimensionality D allows us to treat each of the D GPs as independent. This allows us to compute the total differential entropy at each location of the map by summing the entropy of each independent GP, using the following equation:

$$H(v) = \sum_{d=1}^D \frac{1}{2} \ln[(2\pi e)^D \det(\Sigma_d)] \quad (3)$$

This refers to the entropy at state, v , the GP input vector that contains the spatial location (latitude, longitude) as well as the low-resolution raw spectral data. Σ_d is a single GP covariance matrix of all states in the map computed with a radial basis function (RBF) kernel described in Equation 2.

As the number of points in the map increases, this entropy computation quickly becomes expensive and infeasible, especially on the limited computing capacity of space-grade hardware. Specifically the calculation of the determinant of the predicted covariance matrix, Σ_d . Therefore, we use the upper bound on

the entropy $\hat{H}(v)$ instead, which is more efficiently computed with the main diagonal of the predicted covariance:

$$H(v) \leq \hat{H}(v) = \sum_{d=1}^D \frac{1}{2} \ln[(2\pi e)^D \sigma_d^2] \quad (4)$$

where σ_d^2 is the predicted variance of each point in the map [8]. Therefore the upper bound of the map entropy will be the information metric we will utilize for informative path planning. The entropy at each map location is computed in this way to generate a spatial representation, an entropy map.

3.3 Ergodic Trajectory Planning

Ergodic theory is a mathematical framework that studies the time-averaged behaviour of dynamical systems. A system is said to exhibit ergodic dynamics if it visits every subset of the state space with a probability equal to the measure of that subset, essentially covering a space proportionally to a distribution [45]. The field of ergodic theory historically focuses on analyzing the ergodicity of specific dynamical systems but more recent research focuses on the problem of designing ergodic trajectories, often framed as coverage problems [28], [37], however ergodic trajectory design can be used for IPP by using a measure of expected information in the ergodicity metric.

3.3.1 Ergodicity Metric

In the context of IPP, ergodic trajectory optimization computes controls that drive a dynamic system along trajectories such that the amount of time spent in regions of the state space is proportional to the expected information gain in those regions. This is built upon integrating an ergodicity metric into the objective function to optimize. The ergodicity metric refers to the difference between the time-averaged behavior of the robot trajectory and a spatial prob-

ability density function representing the expected information density (EID).

We will be using the upper bound on entropy from Equation 4 as our EID.

We use an ergodicity metric proposed by Mathew and Mezić [37], which provides distinct benefits balancing exploration and exploitation. The ergodicity metric is the distance of the time-averaged trajectory from being ergodic with respect to the EID, $\phi(x)$. This distance can be quantified by defining a norm on the Fourier coefficients of both distributions. This norm is the sum of the weighted squared distance between the Fourier coefficients of the spatial distribution, ϕ_k , and those of the distribution representing the time-averaged trajectory, $c_k(x(t))$. The ergodicity metric will be defined as $\mathcal{E}(x(t))$, as follows:

$$\mathcal{E}(x(t)) = \sum_{k=0}^K \Lambda_k |c_k(x(t)) - \phi_k|^2 \quad (5)$$

where K is the number of basis functions and Λ_k is a weighting factor which allows for preference of features by frequency due to the nature of the Fourier decomposition. We will be using a Λ_k weighting factor defined by [37].

$$\Lambda_k = \frac{1}{(1 + \|\frac{k\pi}{L}\|^2)^s} \quad (6)$$

Such that $s = \frac{m+1}{2}$, where m is the number of spatial dimensions in the trajectory space. The Fourier basis used in the decomposition are defined as follows:

$$F_k(x(t)) = \frac{1}{h_k} \prod_{x_i \in x(t)} \cos\left(\frac{k_i \pi x_i}{L_i}\right) \quad (7)$$

where the boundaries of each spatial dimension of the distribution are defined from 0 to L_i , k_i is the Fourier basis number, and h_k is a normalization factor defined in [37]. The Fourier decomposition of the time-averaged trajectory, $c(x(t))$, and the EID, $\phi(x)$, are defined below:

$$c_k(x(t)) = \frac{1}{T} \int_0^T F_k(x(t)) dt \quad (8)$$

$$\phi_k = \int_X \phi(x) F_k(x) dx \quad (9)$$

with trajectory planning horizon final time, T .

3.3.2 Ergodic Spectral Multi-scale Coverage

The first of the ergodic trajectory planning algorithms we use is Ergodic Spectral Multi-scale Coverage (SMC) proposed by [37]. This formulation assumes a robotic agent with either first-order or second-order dynamics, we use the first-order dynamics model as defined by:

$$\dot{x}(t) = f(x(t), u(t)) = u(t) \quad (10)$$

The objective is to design a feedback law so that the agent has ergodic dynamics with respect to the supplied EID. The objective function is designed to maximize the rate of decay of the ergodicity metric at the end of a short time horizon. This is referred to as receding horizon control [37], where the control action is determined by solving a finite horizon optimal control problem. The optimization problem leads to a control action that is used at the current time, the beginning of each planning horizon. The computation of the control action is then repeated at every time step. Mathew and Mezić [37] derive the SMC feedback law to be the limit as the size of the receding horizon goes to zero. This provides a closed form solution to the optimization problem and a feedback law efficiently computed at every time step. The objective function to be optimized is as follows:

$$\begin{aligned}
J(t, \Delta t) &= \dot{\mathcal{E}}(x(t + \Delta t)) \\
&= \sum_{k=0}^K \Lambda_k(t + \Delta t) (c_k(x(t + \Delta t)) - \phi_k) (F_k(x(t + \Delta t)) - \phi_k)
\end{aligned} \tag{11}$$

This optimal control problem is solved over the finite time horizon $[t, t + \Delta t]$. As shown in Equation 11, the objective function is the first time-derivative of the ergodicity metric $\mathcal{E}(x(t))$, at the end of the horizon. Note that the controls are subject to the velocity constraint:

$$\|u(t)\|_2 \leq u_{max} \tag{12}$$

Mathew and Mezić [37] solve this optimal control problem in terms of the Lagrange multipliers and the Hamiltonian. The details of which are outlined in [37]. The optimal control solution $u^*(t + \Delta t)$ is the value of an admissible control $u(t + \Delta t)$ that minimizes the Hamiltonian, shown in Equation 13. Specifics on deriving the Hamiltonian and Lagrange multiplier equations when solving optimal control problems are provided in [31].

$$u^*(t + \Delta t) = \underset{\|u(t + \Delta t)\|_2 \leq u_{max}}{\operatorname{arg\,min}} H(x, u, t + \Delta t, c_k, \phi_k, F_k) \tag{13}$$

The feedback control law at each time step, $u^*(t)$, is derived from the solution to this optimization problem and taking the limit as Δt goes to zero.

$$\begin{aligned}
u^*(t) &= -u_{max} \frac{B(t)}{\|B(t)\|_2} \text{ if } B(t) \neq 0, \\
\text{where } B(t) &= \sum_{k=0}^K \Lambda_k t (c_k(x(t)) - \phi_k) \nabla_x F_k(x(t))
\end{aligned} \tag{14}$$




Figure 7: Example Spectral Multi-scale Coverage trajectory on a simulated expected information density. Note that the amount of time spent in regions of the state space is proportional to the expected information density.

The Spectral Multi-scale Coverage (SMC) optimal feedback control law is then computed at each time step to produce ergodic trajectories. The SMC algorithm is outlined below:

An example SMC trajectory is shown in Figure 7 on a simulated EID where lighter color is correlated to greater expected information. Note the inherent multi-scale nature of the trajectory, with large scale, low frequency features being explored first followed by revisiting locations to capture small scale, high frequency features.

Note the SMC algorithm does not inherently consider control effort cost and at each time step moves in the direction to most quickly reduce ergodicity from a given initial position. While SMC effectively covers regions proportional to the EID, the observed behaviour of looping and backtracking is wasteful when considered control cost in addition to information gathering.

3.3.3 Ergodic Projection-based Trajectory Optimization

A second ergodic trajectory planning algorithm we investigate is Ergodic Projection-based Trajectory Optimization (PTO) proposed by [38]. In this formulation, the goal of ergodic trajectory optimization is to solve for the continuous trajectory which minimizes an objective function based on both the ergodic metric and the control effort over a time horizon. Thus PTO differs from SMC in that ergodicity is directly optimized over the entire trajectory, compared to the rate of change of ergodicity over a single step. PTO also allows use to intuitively tune parameters to balance information gain and path length. The PTO objective function is defined as follows:

$$\begin{aligned}
 J(x(t), u(t)) &= q\mathcal{E}(x(t)) + \int_0^T Ru(\tau)^2 d\tau + J_b(x(t)) \\
 \text{s.t. } \dot{x}(t) &= f(x(t), u(t)) \quad x(t_0) = x_0
 \end{aligned}
 \tag{15}$$

where $q \in \mathbb{R}$ and $R \in \mathbb{R}^{m \times m}$ are arbitrary parameters defining the relative importance of minimizing ergodicity vs. control effort. $J_b(x(t))$ is a boundary cost to penalize the trajectory leaving the designated region boundaries. We used a boundary cost defined by [40].

$$J_b(x(t)) = c_b \sum_{n=1}^N \max(x_n - L, 0)^2 + \min(x_n, 0)^2
 \tag{16}$$

where c_b is an arbitrary parameter for weighting the cost of leaving the bounded region. Minimization of the objective function is performed using a projection-based steepest descent optimization method outlined in [25]. This formulation provides distinct advantages over search based methods in that general nonlinear systems can be controlled with dynamics of the following form, without additional layers of controllers.

$$\dot{x}(t) = f(x(t), u(t)) \quad x(t_0) = x_0 \quad (17)$$

For simplicity of comparing trajectory planners, we define our robot dynamics as single integrator first order dynamics, defined as follows:

$$\dot{x}(t) = u(t) \quad (18)$$

Unfortunately, this optimization problem is non-convex and there is no efficient way to find the global optimum. However, we can use an iterative gradient descent method, specifically a projection-based descent [25]. The projection method is attractive because nonlinear constraints are removed during calculation of the descent direction, making the computation much simpler. However, the resulting trajectory might not be dynamically feasible, so it is projected back into a feasible space with a projection operator. The projection operator [25] takes the form of a stabilizing feedback control law which maps any trajectory, whether feasible or infeasible, $(\alpha(t), \mu(t))$ to a feasible trajectory $(x(t), u(t))$.

$$\mathcal{P}(\alpha(t), \mu(t)) = \begin{cases} u(t) = \mu(t) + K(t)(\alpha(t) - x(t)) \\ \dot{x}(t) = f(x(t), u(t)) \end{cases} \quad x(t_0) = x_0 \quad (19)$$

The objective therefore is reformulated as the following, where the trajectory is no longer constrained to the feasible trajectory manifold.

$$J(\mathcal{P}(\alpha(t), \mu(t))) \quad (20)$$

We can now solve for the descent direction, $(z(t), v(t))$, iteratively using a Linear Quadratic (LQ) techniques. The LQ formulation of the descent direction is derived in ([25]), and includes objective function gradients we name a_n and b_n for convenience of notation.

$$\begin{aligned}
& \arg \min_{(\alpha(t), \mu(t))} \int_0^T a(t)^T z(t) + b(t)^T v(t) + z(t)^T Q z(t) + v(t)^T R v(t) dt \\
& s.t. \quad \dot{z}(t) = f(z(t), v(t)) \quad z(t_0) = 0
\end{aligned} \tag{21}$$

We define the gradients need to determine the descent direction.

$$\begin{aligned}
a_n &= \nabla_{x_n} J(x(t), u(t)) \\
&= q \nabla_{x_n} \mathcal{E}(x(t)) + \nabla_{x_n} J_b(x(t)) \\
&= \frac{2q}{T} \sum_{k=0}^K \Lambda_k (c_k(x(t)) - \phi_k) \cdot \nabla_{x_n} F_k(x(t)) + \nabla_{x_n} J_b(x(t))
\end{aligned} \tag{22}$$

For our 2D domain, the gradients are as follows.

$$\nabla_{x_n} F_k(x(t)) = -\frac{\pi}{h_k} \begin{bmatrix} \frac{k_1}{L_1} \sin\left(\frac{k_1 \pi x_1}{L_1}\right) \cos\left(\frac{k_2 \pi x_2}{L_2}\right) \\ \frac{k_2}{L_2} \cos\left(\frac{k_1 \pi x_1}{L_1}\right) \sin\left(\frac{k_2 \pi x_2}{L_2}\right) \end{bmatrix} \tag{23}$$

$$\nabla_{x_n} J_b(x(t)) = 2c_b(\max(x_n - L, 0) + \min(x_n, 0)) \tag{24}$$

The gradient with respect to the control input is as follows.

$$\begin{aligned}
b_n &= \nabla_{u_n} J(x(t), u(t)) \\
&= R u_n \delta t
\end{aligned} \tag{25}$$

Finally the step size, γ , to descend in the descent direction is computed using an Armijo line search [2].

The projection-based ergodic trajectory optimization can be described by Algorithm 1 [38].

An example PTO trajectory is shown in Figure 8 on a simulated EID where lighter color is correlated to greater expected information. Note the inherent multi-scale nature of the trajectory, with large scale, low frequency features

Algorithm 1 Projection-based Ergodic Trajectory Optimization

- 1: Initialize EID with entropy map
 - 2: Initialize trajectory
 - 3: Initialize convergence criteria, ergodicity, directional derivative, max iterations
 - 4: **while** not converged **do**
 - 5: Compute state and control gradients along trajectory
 - 6: Linearize dynamics w.r.t. trajectory
 - 7: Solve Riccati equation and compute descent direction in projection space
 - 8: Compute descent step size with Armijo line search
 - 9: Descend and project the update
 - 10: Compute convergence criteria, ergodicity, directional derivative, iterations
-

being explored first followed by revisiting locations to capture small scale, high frequency features.

Note the PTO algorithm produces trajectories that are typically shorter and involve fewer loops and less backtracking than SMC and other informative path planners. This is expected with the incorporation of control effort into the objective function. Of course, the behavior of PTO is influenced by the relative weighting of hyperparameters, the initialization trajectory, the convergence criteria, number of optimization iterations, and initial position. It is possible to get stuck in a local minimum due to the influence of any combination of these factors. We therefore will explore the sensitivity of these parameters on the performance of the PTO planner.




Figure 8: Example Projection-based Trajectory Optimization trajectory on a simulated expected information density. PTO is initialized with a trajectory computed from SMC shown in blue. Note the dynamics are still highly ergodic but the addition of a control effort cost allows us to better balance information gathering and path length.

4 Simulated Field Experiments

4.1 Data Sets

To perform our experiments we simulate a rover exploration scenario with real spectroscopic data from a region of diverse mineralogy in Cuprite, Nevada [54]. We use high-resolution data to represent *in situ* rover samples and low-resolution data to represent remote data. Specifically, we use the data obtained by Airborne Visible Near Infrared Spectrometer - New Generation (AVIRIS-NG) [22] [23] as proxy for *in situ* spectra. For remote sensing measurements, we use data obtained by the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) [18].

4.1.1 AVIRIS-NG

AVIRIS-NG measures reflectance in the range of $[0.38 \mu\text{m}, 2.5 \mu\text{m}]$ at a spectral resolution of 5nm and a high spatial resolution in the range of $[0.3\text{m}, 4\text{m}]$. Spectra are measured as hyper-spectral images with 600 cross-track elements, providing it with $> 95\%$ cross-track spectral uniformity and $\geq 95\%$ spectral IFOV uniformity. The system also consists of an INS/GPS and an on-board calibrator that performs automated calibration of raw spectra. AVIRIS-NG's navigation data is also used to access surface elevation information from a global topographic dataset [22]. Previous studies have shown that AVIRIS-NG measurements are an appropriate analog to *in situ* infrared spectra [55], with inferences drawn on AVIRIS-NG data transferring well onto data collected by *in situ* sensing instruments. The high spatial resolution also allows us to perform simulated experiments using AVIRIS-NG data as *in situ* spectra. AVIRIS-NG has been used in various science applications such as atmospheric correction [19], ecology and vegetation classification [48], geology and mineralogy [53], as well as coastal bathymetry. [49].

4.1.2 ASTER

ASTER is an imaging spectrometer onboard the Terra spacecraft, a satellite in NASA’s Earth Observation System. The instrument consists of multiple cameras possessing three visible short-wave infrared (VSWIR) and six short-wave infrared (SWIR) bands. Compared to AVIRIS-NG, ASTER has a lower spatial resolution in the range of [15m, 90m]. This resolution is comparable to other spectroscopic instruments such as the Compact Reconnaissance Imaging Spectrometer for Mars (CRISM) on the Mars Reconnaissance Orbiter [42]. The ASTER dataset is a nominal representation of remote data available to guide the exploration of a *in situ* agent.

Several studies have used AVIRIS-NG and ASTER measurements as proxies for low and high resolution spectral measurements, respectively [34], [58], [56]. We evaluate our approach on measurements taken at a mining district in Cuprite, Nevada; a well-studied site with high mineralogical diversity [54]. We were able to associate the two instruments’ observations by aligning them with respect to both their spatial and spectral dimensions. We performed map registration on both images with a planar homography approach. We then used the empirical line method [52] to find the correspondence between the ASTER and AVIRIS-NG reflectance values.

4.2 Spectral Reconstruction Experiments

The goal of these experiments is to evaluate and compare planning methods that generate informative paths that reduce entropy and reduce the error of the reconstructed high-resolution spectra. We now explain the experimental design by which we evaluate the performance of our informative path planners. We first define the test parameters and assumptions, then provide context on the implementation details of the various types of planners to inform experimental

design choices.

4.2.1 Experimental Design

The simulated experiments used to investigate our planners involves using real spectroscopic data from a low-resolution remote instrument and a high-resolution instrument representing the *in situ* data. Each of our path planners relies on our Gaussian Process (GP) regression model to provide entropy information to guide search but also the GP model allows us to reconstruct high-resolution spectra from low-resolution remote data by incorporating information from the collected *in situ* data. This is outlined in detail in Section 2.3. The reconstructed spectra is compared to our ground truth dataset, the dataset representing *in situ* measurements (AVIRIS-NG). The performance of our planners will be evaluated by the mean final entropy of the site and the mean final root mean square reconstruction error (RMSE). In short, lower entropy and lower RMSE is better performance. The results are dependent on GP model parameters however the GP parameters are hand tuned prior to experiments on a withheld portion of the dataset.

Average values for each experimental site are reported as well as statistical analysis to determine if planner performance is statistically different with 95% confidence. We use five unique sites from the Cuprite datasets. Figure 9 shows the larger Cuprite region and the five subsections that are our experimental sites. Sites A, B, and C were chosen for their mineralogical diversity and spectral diversity. Site D was chosen for its lack of spectral diversity with large homogeneous regions, providing insight into how the planners balance exploration and exploitation. Finally, Site E was chosen to be a much larger region to give insight into how the planners perform with longer trajectories. All planners are run on all sites from the same 36 initial positions. 36 uniformly spaced initial positions are defined, deterministic planners are run once for each initial

position while stochastic planners are run 5 times for each initial position with the average of the 5 scores reported for that initial position.

4.2.2 Baseline Planners

We compare a variety of path planners’ ability to generate informative paths to minimize total entropy of a region. We use a number of baseline comparison planners in addition to the ergodic planners. Our comparison planners include a random walk planner (Random), a greedy entropy planner (Greedy), an information gain based Monte Carlo tree search planner (MCTS).

- Random: In this case, the sampling location is randomly selected within a fixed radius of the robot’s current location. The next location is sampled from a uniform distribution and does not require prior discretization of the state space. This planner is uninformed, with no notion of information or entropy and acts as a lower bound on planner performance.
- Greedy: The greedy entropy planner selects sampling locations that contain the maximum entropy within a fixed radius from the current robot location. The entropy at each map location is updated using a Gaussian Process update step with each new sample taken and before the next sample location is computed. This planner also does not rely of discretization of state space due to the continuous nature of Gaussian Processes. This planner is functionally similar to the concept of Maximum Entropy Sampling [50] but includes the constraint of a maximum step size from the current location. This planner provides a better low bound comparison with an informative planner, although a naive simple planner.
- MCTS: The information gain based Monte Carlo tree search planner was developed by Kodgule et al. and is decribed in detail in [32]. This planner uses a Markov Decision Process formulation and a Monte Carlo Tree




Figure 9: The boundaries for each experimental site in the Cuprite, Nevada region. For site A we plan 10 samples, for sites B, C, and D we plan 20 samples, and for site E we plan 50 samples.

Search method as a solver. This approach integrates closely with our GP model with a information gain metric based objective function. High-resolution spectra is simulated by the planner and the change in entropy is computed as an information gain metric. The planner uses a look-ahead depth parameter we keep fixed at 5 steps. The planner will expand branches to simulate the total information gain after 5 steps and return the first step of the branch with the maximum information gain. With each execution of a single step the Gaussian Processes are updated with the new sample before a new step is planned. The formulation of the MCTS planner does rely of discretization of the state space, the planner utilizes an 8-connected grid who resolution is defined by the hyperspectral image. It is important to note, MCTS has no conception of its total path length during execution, just a look-ahead depth. Computation time is also often orders of magnitude larger than other planners. Therefore, MCTS is non-myopic and provides a high performance bound.

4.2.3 Ergodic Planners

We are interested in interrogating the performance of our various ergodic informative path planners compared to the baseline planners. We explore two types of ergodic path path planners, Spectral Multi-scale Coverage (SMC) and Projection-based Trajectory Optimization (PTO). We describe in detail the mathematical foundation of these ergodic planners in Section 3.3.2 and Section 3.3.3 for SMC and PTO respectively. We provide a higher level summary and provide insights into implementation details to better express their distinction from the baseline planners.

- SMC: Spectral Multi-scale Coverage is an implementation of a closed form solution to the optimal control problem outlined in Section 3.3.2. The

objective function to be optimized is the rate of decay of the ergodicity metric described in Section 3.3.1. The solution in [37] is only derived for first-order and second-order dynamic systems. We implement the first-order system dynamics to plan paths as the other baseline planners do not require a dynamics model. SMC plans one sample at a time but can return a full trajectory as well, given that ergodic planners contain an internal belief state, the Fourier decomposition of the provided entropy map, $\phi_k(x)$. The only parameter is the maximum step size, which is defined to match the baseline planners. The time averaged statistics of the robot state, $c_k(x(t))$, are compared against $\phi_k(x)$ to compute the next optimal sample locations.

- PTO: Projection-based Trajectory Optimization, proposed by Miller [38], optimizes over the entire trajectory using an objective function based on both the ergodic metric and the control effort. This allows us to intuitively tune the parameters q and R , the relative importance of reducing map entropy and path length respectively. The mathematical details of PTO are described in Section 3.3.3. The weighting parameters are hand tuned to roughly estimate the average step size of the other planners for a fair comparison. PTO optimizes over the entire trajectory, in other words is only minimizing the final entropy of the scene once all samples in the trajectory have been collected. Similar to the MCTS look ahead feature, this allows the planner to more intelligently plan. For example, if the planner has a conception of its planning horizon it may find a path is more optimal if it moves through a less informative region to reach a more informative region. This provides an advantage over myopic planners but even non-myopic planners, like MCTS, which still require a shorter look ahead to be computationally feasible [32]. PTO uses an iterative

projection-based optimization approach that utilizes a gradient method for optimization and therefore requires an initialization trajectory. We explore initializing the PTO planner with and SMC trajectory and with a random walk in cardinal directions. PTO also guarantees a dynamically feasible trajectory for the given dynamics model but has the benefit of being suitable for arbitrary non-linear dynamics.

It is important to note that the simulated field experiments are able to be run without being constrained to a discretized grid by use of bilinear interpolation [24]. The Random, Greedy, SMC, and PTO planners plan sub-pixel sample locations that are approximated by bilinear interpolation. This approximates a linear spectral mixture model [26].

Table 2: Planner Property Comparison

Planner Property	Random	Greedy	MCTS	SMC	PTO
Fixed to grid			✓		
Smooth paths				✓	✓
Optimal over trajectory					✓
Plans single step	✓	✓	✓	✓	
Plans whole trajectory				✓	✓
Dynamically feasible paths				✓	✓
Non-linear dynamics					✓

4.2.4 Planning Horizon - Model Predictive Control

Each of the informative baseline planners Greedy and MCTS performs a Gaussian Process (GP) update after each sample to then plan the next sample location with an updated entropy map. Entropy across the scene will reduce with each sample given the spatial and spectral correlation between of input data to the GPs. This process is described in detail in Section 2.3. The ergodic planners due to their internal belief state can plan whole trajectories using only the initial

entropy map, and these trajectories are often competitive with high performance baseline MCTS. It follows that the ergodic planners should be able to improve performance by using a Model Predictive Control (MPC) framework and iteratively re-planning with shorter planning horizons. The ergodic planners can plan a full trajectory of N samples from an initial entropy map, n steps of the trajectory will be executed and the entropy map updated before re-planning the remaining trajectory. This can be performed with arbitrary planning horizon n . We therefore experiment with planning horizon and examine its effects on ergodic planner performance. For all ergodic planners including SMC, PTO with SMC initialization, and PTO with random initialization the planning horizon is set to re-plan after 1 sample, 4 samples, 8 samples, and to plan the entire trajectory only with the initial entropy map (no re-planning). For clarity, the naming convention of the planners includes "mpc" and a planning horizon (i.e. smc-mpc-1 re-plans after every 1 sample) when iterative planning is used. When the "mpc" is excluded, this indicates no re-planning is performed.

5 Results and Analysis

5.1 Planner Performance Overview

We evaluate the performance of the informative path planners by the mean final mean entropy of the site and the mean final root mean square reconstruction error (RMSE) of the map. In short, lower entropy and lower RMSE means better performance. The mean entropy of all map locations is reported after each sample is updated into the Gaussian Process (GP) spectroscopic model, computed by Equation 4. The RMSE is computed by predicting spectra from the GP spectroscopic model at every map location where there is *in situ* analog data (AVIRIS-NG for Cuprite) and comparing against the real data as a ground truth. Ultimately the entropy metric is more indicative of planner performance as it is the metric that is directly being optimized by the planners, RMSE is more sensitive to the tuning of the GP model. RMSE results show a high correlation between entropy reduction and accuracy of reproducing high-resolution spectra [8]. We compute an efficiency metric meant to express the average efficiency of each planner to reduce entropy per a unit of path length, higher efficiency is better. We also compute the average total computation time for planning each trajectory and performing the GP model updates.

Table 3: Best Performing Planner by Site

Site	Final Entropy	Final RMSE	Final Entropy Efficiency	Computation Time
A	pto-rand-init-mpc-1	pto-rand-init-mpc-1	random	smc
B	pto-smc-init-mpc-1	pto-rand-init-mpc-1	pto-rand-init	pto-smc-init
C	pto-smc-init-mpc-1	pto-smc-init-mpc-1	smc-mpc-1	random
D	pto-smc-init-mpc-1	pto-smc-init-mpc-1	pto-smc-init	random
E	pto-smc-init-mpc-1	pto-smc-init-mpc-1	pto-rand-init-mpc-8	smc

Best performing planner for each site for each performance criteria.

The best performing informative path planner for each evaluation criteria is shown in Table 3. PTO is shown to consistently perform best in terms of entropy and RMSE. Although PTO-SMC-INIT-MPC-1 (PTO with SMC initialization) and PTO-RAND-INIT-MPC-1 (PTO with random initialization) are shown to be the best performing planners for entropy and RMSE, their performance difference is not statistically significant with 95% across all sites. This indicates that either PTO initialization is the highest performing planner. However the planning horizon of 1 sample clearly outperforms longer planning horizons across all sites, and the performance difference between 1 sample planning horizon and others is very often statistically significant.

The entropy efficiency varies with planning horizon, algorithm, and initialization. Even Random performs best for site A, largely due to the small size of the site and the short path lengths that often result from Random plans. On average ergodic planners perform with the highest efficiency. There is also no clear best planner for computation time but on average the closed form planners performed quicker than those that require a solver or optimizer. We can decisively state that MCTS is the slowest planner, consistently taking two orders of magnitude longer than others.

On average Random performs poorly but provides a good low baseline that shows entropy is reduced even when taking uninformed samples with a short trajectory. The Greedy planner is highly susceptible to noise and often gets stuck in sub-optimal minima. MCTS often produces rectilinear paths due to being fixed to a grid. The paths also often cross back over themselves. All ergodic planners produce smoother paths than non-ergodic planners. SMC often is competitive with MCTS despite SMC only planning single steps and MCTS having a look-ahead depth of 5 samples. For some sites SMC even outperforms MCTS in terms of entropy and RMSE. Note that Figure 11 shows that both

PTO planners only perform best after nearly all samples are collected. This behavior is expected as the planners optimize over the whole trajectory not each point individually.

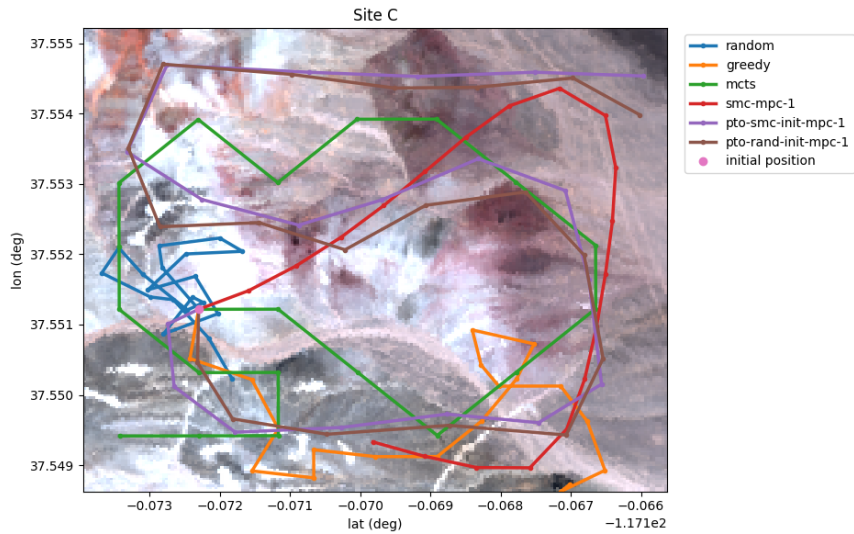


Figure 10: This example run shows typical path shapes that each planner achieves. Note the smooth paths of the ergodic planners that efficiently cover the region and the non-smooth MCTS path which backtracks. Random produces short, jagged, uninformed paths while Greedy often gets stuck in a local minimum somewhere on the map.

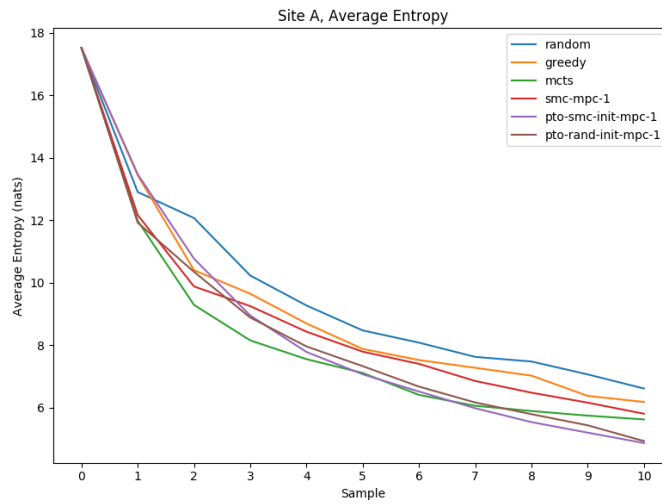


Figure 11: Average of all runs reported, lower entropy is better performance. Random performs the worst and performance improves as the planners become less myopic. PTO performs 31.2% better than Random and 16.3% better than Greedy in terms of entropy reduction. The two PTO initialization's performance is not statistically significantly different.

5.2 Ergodic Planners

Ergodic planners in general provide particular advantages over the other baseline planners. They notably do not require prior discretization of the state space, like MCTS, yet consistently perform as well or better than MCTS. This greatly improves the ability of the planner to collect informative samples independent of spatial frequency. The computation times required are often on par with Random and Greedy and are significantly faster than MCTS. The ergodic planners also generate smooth paths that are guaranteed to be dynamically feasible based on the dynamics model used internally to the planner. This makes them a strong candidate to be used with underactuated robots like AUVs and UAVs. Shown in Figure 12 are example paths from site E, the ergodic planner path is significantly smoother and therefore energy efficient for a number of robot systems.

The PTO algorithm consistently performs the best in terms of entropy reduction and RMSE. This is expected as it is highly non-myopic, optimizing over the entire trajectory. This also explains why in figures like Figure 11, the entropy of PTO does not outperform all other planners until nearly all samples have been collected. The optimization is performed with a projection-based gradient method and therefore can place the sample locations optimally for a given EID often without backtracking. The parameters allow the user to intuitively balance the relative importance of minimizing ergodicity and control effort. This allows for path length to be controlled indirectly.

The SMC algorithm displays multi-scale behavior driven by our formulation of the ergodicity metric in Equation 5. The weighting vector in Equation 6 shows larger weighting for low-frequency than high-frequency features. Therefore SMC trajectories tend to first visit locations with low-frequency broad information features then visit high-frequency features. The paths therefore are often smooth

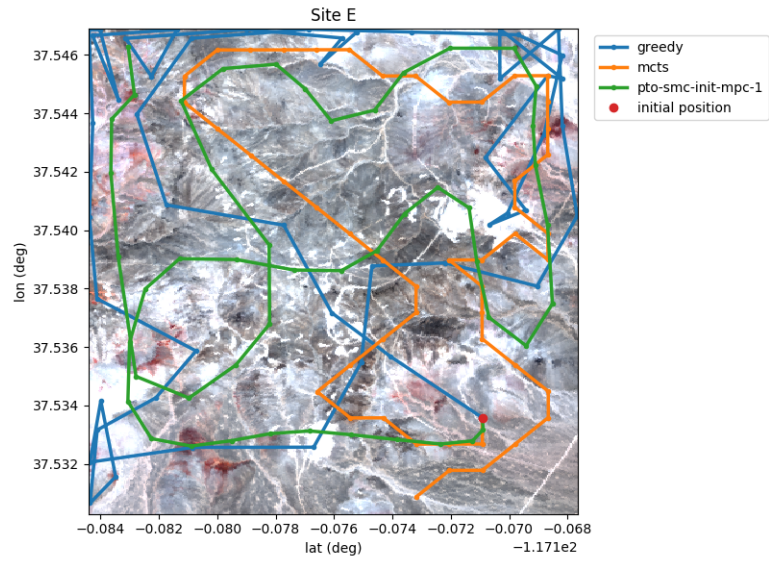


Figure 12: Ergodic planners produce smooth paths compared to baselines and therefore are good candidates for use with AUVs and UAVs.

and become less smooth as the trajectories get longer. This is shown in Figure 13.




Figure 13: Example SMC trajectory displaying multi-scale behavior. The path first visits locations with low-frequency information features and then prioritizes high-frequency features.

5.3 Variation Between Sites

We report the final performance metrics for all planner types, initialization types, and planning horizons for all sites below. We analyze the performance of the planners in terms of the differences between site sizes, spectral diversity, and spectral uniformity.

The sites vary by size as shown in Figure 9. Sites A, B, and C were selected for their spectral and geologic diversity. They contain numerous mineral classes with distinct spectral signatures and are strong candidates for evaluating planner performance under nominal conditions. Site D was selected to examine the performance of planners on a site with low spectral diversity and large areas of homogeneity. This allows us to examine the planner performance with a less informative prior. Site E was chosen to be very large with numerous areas of spectral diversity to investigate the effect of region size on planner performance.

Table 4: Site A Planner Comparison

Planner	Final Entropy (nats)	Final RMSE	Final Entropy Efficiency	Computation Time (sec)
random	6.921	0.2278	35.23	1.33
greedy	5.951	0.2084	33.65	1.39
mcts	5.618	0.1799	16.27	134.42
smc	5.504	0.1709	31.49	1.27
smc-mpc-8	5.497	0.1717	31.45	1.42
smc-mpc-4	5.542	0.1718	31.22	1.32
smc-mpc-1	5.515	0.1693	31.25	1.48
pto-smc-init	5.329	0.1577	27.31	2.51
pto-smc-init-mpc-8	5.329	0.1562	26.48	3.55
pto-smc-init-mpc-4	5.204	0.1522	25.35	4.66
pto-smc-init-mpc-1	5.088	0.1476	23.81	12.08
pto-rand-init	5.224	0.1511	29.73	2.49
pto-rand-init-mpc-8	5.201	0.1506	29.89	3.67
pto-rand-init-mpc-4	5.245	0.1518	27.84	4.84
pto-rand-init-mpc-1	5.053	0.1468	25.97	11.94

Site A, 10 samples, best performance in bold, final values after all samples are collected for entropy, RMSE reconstruction error, computation time are averaged, entropy efficiency is a measure of entropy vs. path length

On average PTO variants perform best across all sites. On average, Random performs the worst, followed by Greedy, SMC, and MCTS. The results are consistent across sites A, B, and C with the only notable difference is that in site C the MCTS planner is much more competitive with PTO. Site D shows greater spread between the entropy performance of all planners, which is expected based on the uniformity of the region. SMC even outperforms MCTS for site D which gives credence to the claim that ergodic planners are well suited to balance exploration and exploitation. Site E performance also matches the greater performance trend well with SMC notably performing nearly as well the PTO variants.

Table 5: Site B Planner Comparison

Planner	Final Entropy (nats)	Final RMSE	Final Entropy Efficiency	Computation Time (sec)
random	7.672	0.2574	7.05	16.22
greedy	6.191	0.2068	7.49	18.76
mcts	5.498	0.1833	7.53	1112.66
smc	5.664	0.1858	6.70	16.04
smc-mpc-8	5.611	0.1848	6.75	19.22
smc-mpc-4	5.604	0.1842	6.75	17.74
smc-mpc-1	5.591	0.1845	6.76	19.61
pto-smc-init	5.621	0.1854	8.97	15.66
pto-smc-init-mpc-8	5.499	0.1798	8.02	20.44
pto-smc-init-mpc-4	5.519	0.1819	7.30	24.15
pto-smc-init-mpc-1	5.425	0.1807	6.60	51.42
pto-rand-init	5.647	0.1859	9.38	17.12
pto-rand-init-mpc-8	5.569	0.1854	8.44	21.70
pto-rand-init-mpc-4	5.538	0.1843	8.21	24.76
pto-rand-init-mpc-1	5.458	0.1795	6.98	52.57

Site B, 20 samples, best performance in bold, final values after all samples are collected for entropy, RMSE reconstruction error, computation time are averaged, entropy efficiency is a measure of entropy vs. path length

Table 6: Site C Planner Comparison

Planner	Final Entropy (nats)	Final RMSE	Final Entropy Efficiency	Computation Time (sec)
random	8.797	0.2921	10.53	14.46
greedy	6.349	0.1788	12.13	14.53
mcts	4.986	0.1416	8.21	1094.60
smc	5.583	0.1578	12.43	15.05
smc-mpc-8	5.564	0.1589	12.47	17.09
smc-mpc-4	5.537	0.1571	12.53	15.66
smc-mpc-1	5.532	0.1569	12.55	17.86
pto-smc-init	5.162	0.1444	11.23	15.71
pto-smc-init-mpc-8	5.112	0.1393	9.28	20.52
pto-smc-init-mpc-4	5.047	0.1429	8.50	24.41
pto-smc-init-mpc-1	4.970	0.1387	7.66	52.16
pto-rand-init	5.157	0.1470	10.56	18.20
pto-rand-init-mpc-8	5.168	0.1467	10.06	21.42
pto-rand-init-mpc-4	5.142	0.1468	9.70	25.90
pto-rand-init-mpc-1	4.995	0.1410	8.60	54.23

Site C, 20 samples, best performance in bold, final values after all samples are collected for entropy, RMSE reconstruction error, computation time are averaged, entropy efficiency is a measure of entropy vs. path length

Table 7: Site D Planner Comparison

Planner	Final Entropy (nats)	Final RMSE	Final Entropy Efficiency	Computation Time (sec)
random	9.547	0.2978	6.36	27.95
greedy	6.634	0.1856	7.21	32.77
mcts	5.141	0.1434	7.71	2269.47
smc	5.222	0.1452	8.30	31.94
smc-mpc-8	5.231	0.1454	8.32	38.60
smc-mpc-4	5.221	0.1444	8.34	34.87
smc-mpc-1	5.170	0.1423	8.39	40.41
pto-smc-init	5.138	0.1393	8.51	28.44
pto-smc-init-mpc-8	4.728	0.1334	7.51	36.50
pto-smc-init-mpc-4	4.686	0.1301	7.07	38.95
pto-smc-init-mpc-1	4.451	0.1229	6.33	71.80
pto-rand-init	4.712	0.1305	8.11	31.57
pto-rand-init-mpc-8	4.607	0.1267	7.34	36.00
pto-rand-init-mpc-4	4.556	0.1243	7.12	38.13
pto-rand-init-mpc-1	4.501	0.1237	7.07	69.97

Site D, 20 samples, best performance in bold, final values after all samples are collected for entropy, RMSE reconstruction error, computation time are averaged, entropy efficiency is a measure of entropy vs. path length

Table 8: Site E Planner Comparison

Planner	Final Entropy (nats)	Final RMSE	Final Entropy Efficiency	Computation Time (sec)
random	6.388	0.2082	1.30	154.59
greedy	5.077	0.1829	1.48	174.65
mcts	4.910	0.1686	3.21	27891.40
smc	4.786	0.1616	2.38	137.21
smc-mpc-8	4.714	0.1586	2.41	148.28
smc-mpc-4	4.682	0.1597	2.43	156.85
smc-mpc-1	4.702	0.1587	2.42	171.55
pto-smc-init	4.820	0.1630	3.86	138.99
pto-smc-init-mpc-8	4.895	0.1667	3.61	185.94
pto-smc-init-mpc-4	4.792	0.1649	3.28	206.57
pto-smc-init-mpc-1	4.607	0.1570	2.56	386.06
pto-rand-init	4.819	0.1651	3.58	144.88
pto-rand-init-mpc-8	5.207	0.1758	3.97	197.84
pto-rand-init-mpc-4	4.881	0.1665	3.51	216.53
pto-rand-init-mpc-1	4.632	0.1580	2.57	407.35

Site E, 50 samples, best performance in bold, final values after all samples are collected for entropy, RMSE reconstruction error, computation time are averaged, entropy efficiency is a measure of entropy vs. path length

5.4 Planning Horizon

The SMC, PTO-SMC-INIT, and PTO-RAND-INIT planner variants were all run with various planning horizons to investigate changes to performance. Across all sites there is a clear trend that smaller planning horizons results in improved performance. Figure 14 shows the final entropy result statistics for all planners. The median value clearly decreases with planning horizon. This is consistent across all sites as shown in the Appendix.

We also must determine if this difference in performance is statistically significant. The statistical significance does vary more by site than median entropy performance but we will report general trends. SMC variants are more often not statistically different as shown in Figure 15. Dark cells indicate the performance difference is statistically significant with 95% confidence. For the PTO-SMC-INIT and PTO-RAND-INIT variants the 1 sample planning horizon is the only variant that is consistently statistically different than the other planning horizons. This result makes sense because as the planning horizon decreases, the planner has access to more information. When the entropy map is updated more frequently the planners have the most up to date information.

Therefore we conclude updating the entropy map after each sample, a planning horizon of 1 sample, will achieve the best performance. Computation time does increase as a result but the time scales of the planning compared to execution is often negligible. A planning horizon of 4 samples is a favorable balance between information gain and computation time, where in some cases the 4 sample and 1 sample horizon performance is not statistically different.

We will also examine planning horizon effects qualitatively. The paths are less smooth with shorter planning horizons but the detours are intuitively energy efficient. With more frequency entropy map updates more local regions of high entropy are exposed and the planner can adjust accordingly and take detours




Figure 14: The statistical variation of final entropy between runs is shown here. There is a clear trend in median final entropy reducing as planning horizon reduces.

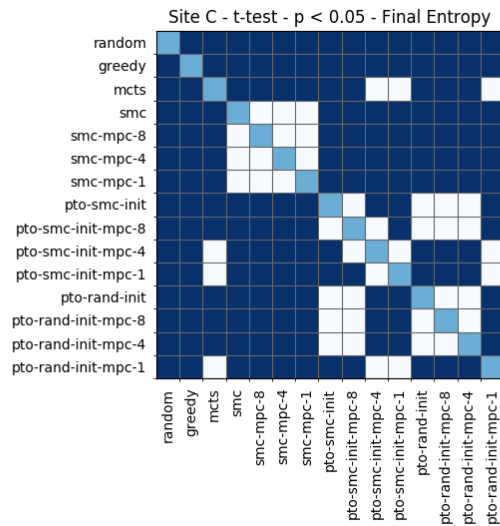


Figure 15: 2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence.

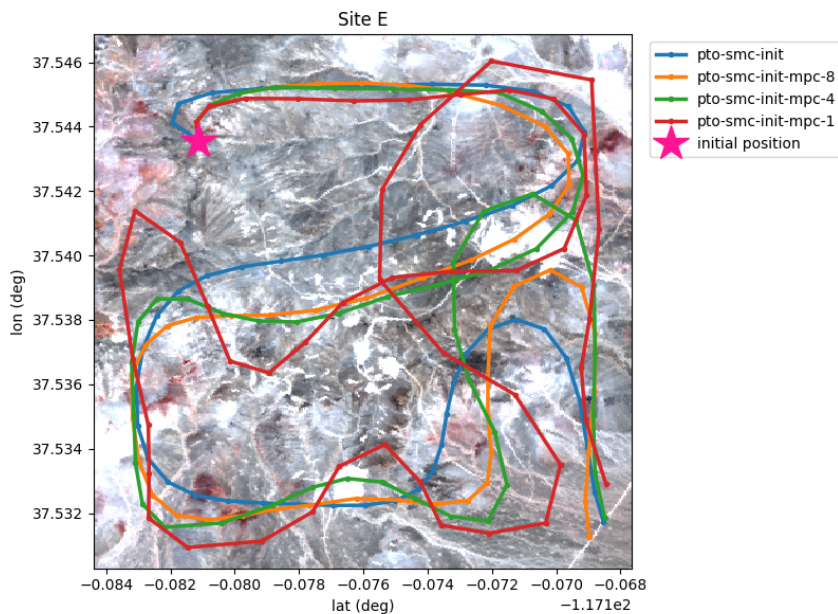


Figure 16: Paths of multiple planning horizons from the same initial position. Notice the paths diverge from each other as the plans execute and as the planning horizon is shortened. The paths get less smooth and have more turns with the shorter horizons as the planners are seeing more localized bright entropy spots with more entropy map updates.

from the previously planned actions. We see this behavior in Figure 16, which also makes sense considering shorter planning horizons perform better in terms of entropy and RMSE.

5.5 Summary

Informative path planning for science autonomy applications has been studied in a limited context. Traditional planners often did not fully utilize prior information available and continue to adaptively use new information. This thesis presented two ergodic informative path planners, Spectral Multi-scale coverage (SMC) and Projection-based Trajectory Optimization (PTO), as viable planning methods for active spectroscopic mapping.

Simulated field experiments were performed across regions of mineralogical and spectral diversity, comparing ergodic path planning against various baseline planners, random, greedy, and MCTS. Over 7000 simulations were performed across 5 sites. The experiments show that PTO consistently outperforms all other planning algorithms in entropy reduction and RMSE with a statistically significant performance difference with 95% confidence. Ergodic planners perform with the best entropy efficiency in 4 out of 5 sites and the computation time required is orders of magnitude lower than the next best baseline, MCTS.

We also investigated the planning horizon used for ergodic trajectory optimization. We performed the simulated field experiments with no re-planning, and re-planning with an updated entropy map after every 8 samples, 4 samples, and 1 sample. Consistently the shortest planning horizon (1 sample) performed best in entropy reduction and RMSE. For some sites, no re-planning and longer planning horizons resulted in better entropy efficiency which suggests planning horizon should be more thoroughly considered for efficiency critical applications.

Ergodic trajectory optimization ultimately is a strong candidate for science autonomy applications. The SMC and PTO variants produce smoother paths that would be ideal for AUV or UAV applications. They do not require discretization of the state space and are computationally efficient compared to baselines. If the total path length is known, PTO is shown to be the best performing planning. In applications where total number of samples is unknown and compute is limited, SMC provides high performance where PTO is infeasible.

6 Conclusion

Robotic explorers will require an increasing degree on autonomy to improve exploration efficiency and mitigate challenges like communication bandwidth and delay. We aim to improve on robotic explorers' ability to plan for scientific discovery and automatically interpret measurements in the field. By encoding belief of the spatial distribution of the information content of a region, we can iteratively plan actions that will guide the robot toward more information rich areas. Our context driven exploration strategy draws inspiration from and captures the characteristics of geologic site survey. Sample interpretations are sensitive to spatial relations. Rather than considering each sample independently, field scientists incorporate each new observation into an evolving model of geologic formation.

We compared the performance of ergodic planning techniques with baseline informative search techniques in terms of entropy reduction, RMSE, computation time, and the sample efficiency of each trajectory. We have demonstrated that the ergodicity metric that defines the objective functions of ergodic planners favorably balances exploration and exploitation of prior information. We show competitive performance of the SMC and PTO planners even when using little information and planning from an initial entropy map. We show that PTO consistently outperforms all other planners in our experiments with a planning horizon of 1 sample. We also conclude that re-planning with updated entropy maps consistently improves performance with an optimal planning horizon of 1 sample.

6.1 Contributions

This thesis contributes to the field of planetary exploration in the following ways:

- Developed novel formulations for active spectroscopic mapping for use with the Ergodic Spectral Multi-scale Coverage (SMC) algorithm and Ergodic Projection-based Trajectory Optimization (PTO) algorithm.
- Designed and implemented a code base for further study using ergodic planning in Python
- Produced datasets and test environment for simulated field experiments
- Validated the efficacy of ergodic trajectory optimization as a tool for informative path planning compared state-of-the-art path planners.

6.2 Future Work

This research set up a foundation for applying ergodic trajectory optimization to science autonomy, but there are many more questions to answer. Extending ergodic trajectory optimization with terrain aware planning that allows robots to avoid obstacles or locations of untraversable slopes would make for more realistic and usable planning for ground robots.

Another research question would be exploring ergodic trajectory optimization’s usefulness as a tool for continuous sampling. Camera sensors and smaller spectrometer instruments allow robots to efficiently sample and analyze large amounts of data on the fly. Because SMC and PTO can extend efficiently to large numbers of sample points they effectively plan a continuous sample path that is smooth.

Finally ergodic trajectory optimization’s practicality as a planning method for UAVs suggests its ability to extend to planning ergodic trajectories in three dimensions for information gathering. Leveraging an aerial robot’s ability to change altitude to efficiently gather information at different scales should improve entropy efficiency overall. Change in altitude allows the robot to quickly

change the field of view and subsequently the information content measured by a spectrometer instrument. The primary issue is the state space extends to a third physical dimension yet the information space is still in two physical dimensions. The ergodic path planning algorithms we used here require the information space dimensions to match the state space dimensions. Heuristic functions and sensor models may provide a method for computing gradients in third dimension to better plan trajectories using PTO. More investigation into the sensor model and how change of field of view and distance to target changes the resultant signal will be required to develop this.

The question of updating the Gaussian Process model in our specific case is also an open question. Adaptively changing the length scale kernel parameter of the Gaussian Process model with altitude or signal noise may provide a solution to updating the model with varying degrees of influence. Another potential solution would be incorporating a non-stationary kernel function which is a function of the full robot state.

A Appendix

A.1 Site A

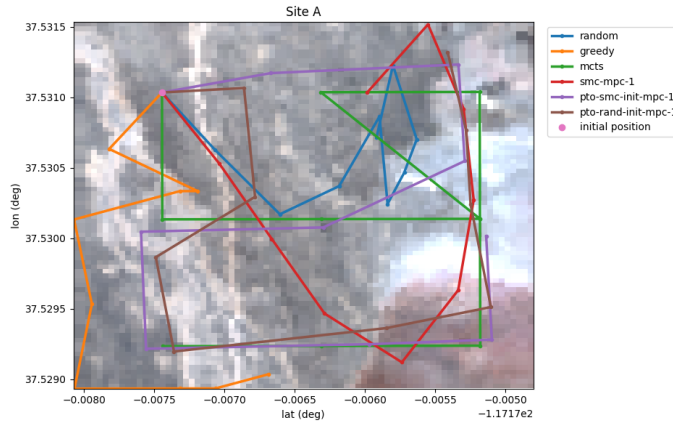


Figure 17: Greedy, as expected, is more susceptible to noise and gets stuck in a very sub-optimal minimum. All ergodic planners produce smoother paths compared to non-ergodic. PTO and MCTS globally look to cover similar points but the longer planning horizon of PTO allows it to perform best. Note the irregular shape of the MCTS path due to being fixed to the grid, MCTS also backtracks and repeats a sample location.




Figure 18: Examples like this show why Random often reports a high efficiency; even though the entropy score is always the worst, the paths are typically much shorter. MCTS again has a non-smooth path that crosses back over itself. Note that PTO-SMC and PTO-RAND move in different directions with different path shapes but consistently outperform baselines.




Figure 19: Average of all runs reported, lower entropy is better performance. Random performs the worst and performance improves as the planners become less myopic. SMC is competitive with MCTS despite SMC only planning single steps and MCTS has a look-ahead depth of 5 samples. Note that both PTO planners only perform best after nearly all samples are collected, this is expected as the planners optimize over the whole trajectory not each point individually.

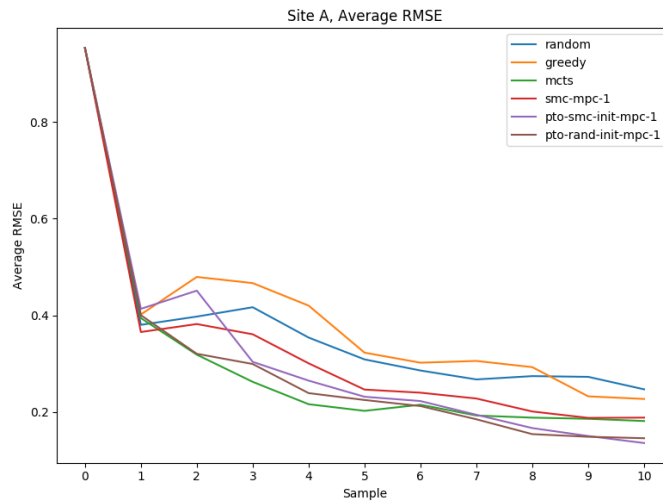


Figure 20: Average of all runs reported, lower RMSE is better performance. Note that signal trends are highly correlated between entropy and RMSE, therefore our entropy calculation is a reliable information metric for use in informative path planning. Notably, the Greedy planner RMSE curve differs greatest from the entropy curve. It is the most myopic of the informative planners and susceptible to noise, i.e. it greatly favors exploitation over exploration. This suggests there that while entropy and RMSE are highly correlated there is noise on the entropy signal that must be accounted for.




Figure 21: The statistical variation of final entropy between runs is shown here. Greedy and SMC have similarly large IQR, this is expected given the single step planning of each. SMC is ultimately able to perform better than Greedy because the ergodicity metric it optimizes over is better able to balance exploration and exploitation. The planning horizon changes between planners does not appear to greatly effect statistical spread of planner performance.

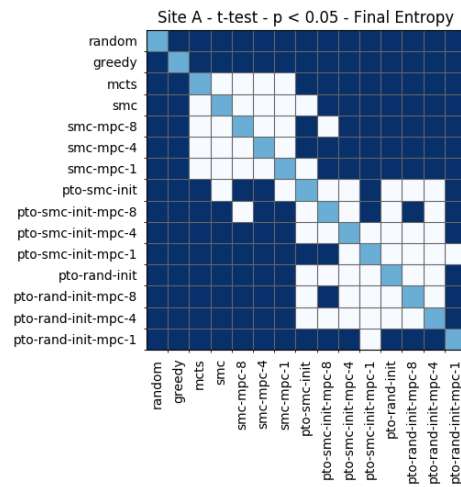


Figure 22: 2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence. MCTS and SMC variant performance are not different. Notably, most PTO variants are different than all other planners, while most PTO variants are not different than each other. PTO-RAND-INIT-MPC-1 performs best and is statistically different almost all other planners.

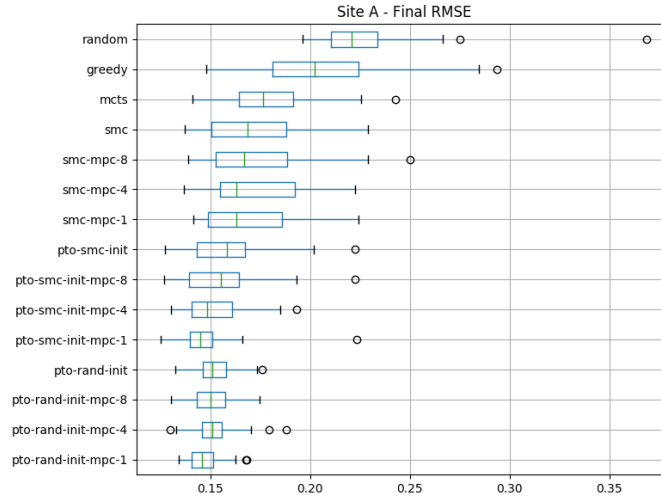


Figure 23: The statistical variation of final RMSE between runs is shown here. Once again, trends of the entropy performance are matched in the RMSE performance. The Greedy planner large IQR once again reflects the planner's susceptibility to noise.

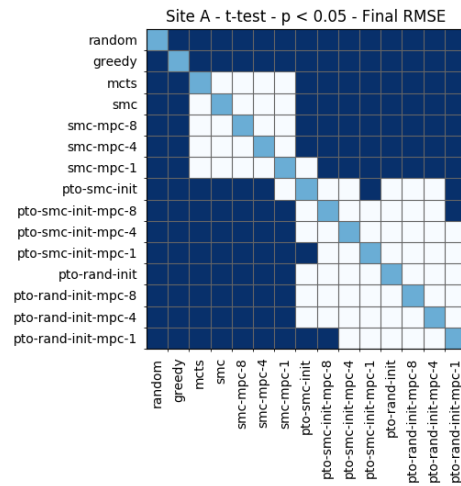


Figure 24: 2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final RMSE performance between planners with 95% confidence. The entropy performance again is reflected. The notable difference is that PTO-RAND-INIT-MPC-1 RMSE performance is not different than most other PTO variants.

A.2 Site B

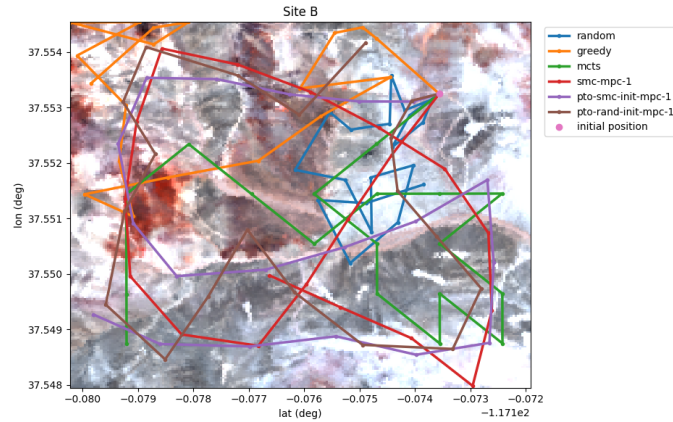


Figure 25: Greedy again is stuck in a local minimum. The multi-scale nature of SMC is shown here, the path targets low frequency information features first before revisiting areas for higher frequency features.

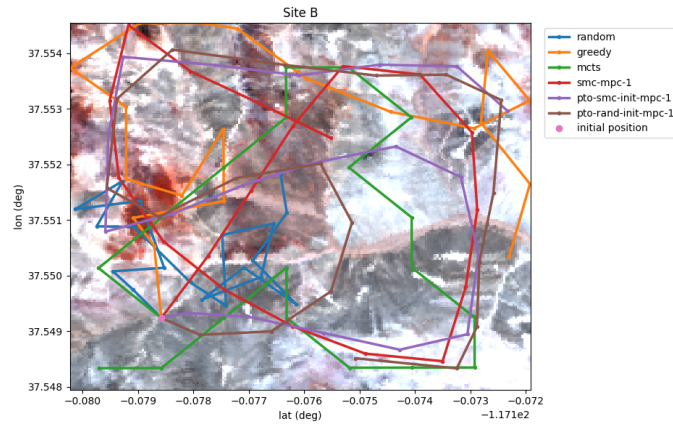


Figure 26: Note how PTO is able to achieve nearly all the same points as SMC but without backtracking due to optimizing over the full trajectory. Given the current science goals moving on to a new site may or may not be desired, however SMC is more likely to return to similar areas given its multi-scale nature.

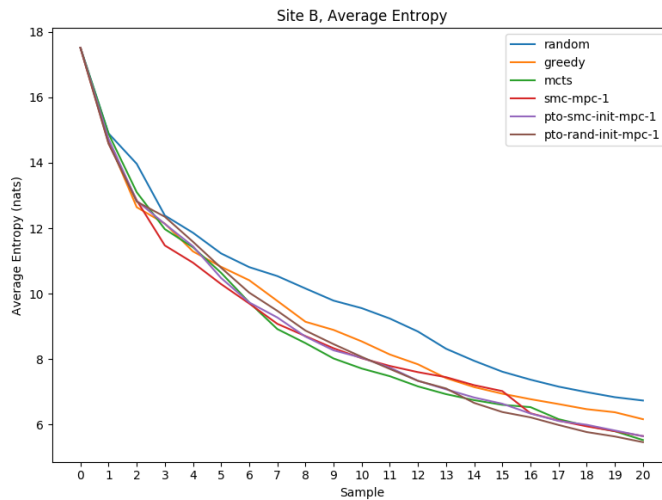


Figure 27: Average of all runs reported, lower entropy is better performance. Random performs the worst and performance improves as the planners become less myopic. We once again see the PTO planners over taking others as the samples approach the full trajectory due to their being highly non-myopic.




Figure 28: Average of all runs reported, lower RMSE is better performance. Note that signal trends are highly correlated between entropy and RMSE, therefore our entropy calculation is a reliable information metric for use in informative path planning.

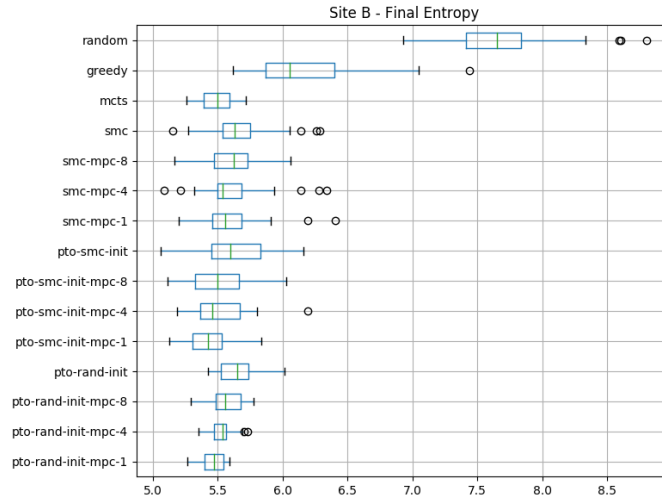


Figure 29: The statistical variation of final entropy between runs is shown here. Greedy has a particularly large IQR indicating there are local minima that myopic planners are susceptible to. SMC’s multi-scale nature allows it to avoid the noisy local minima. There is a clear trend in median final entropy reducing as planning horizon reduces.

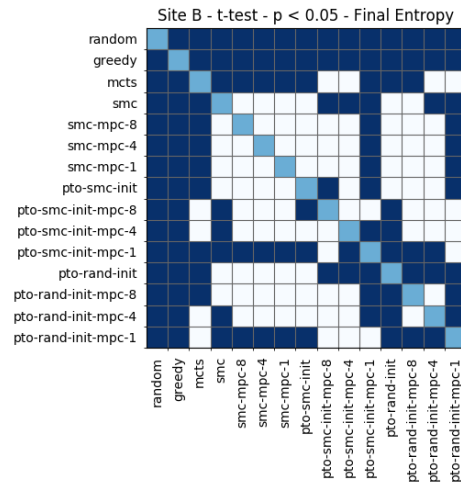


Figure 30: 2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence. MCTS is not statistically different than a few ergodic variants but is notably different than the best performing planner, PTO-SMC-INIT-MPC-1. The SMC variants are not statistically different from each other as are many of the PTO-SMC variants.

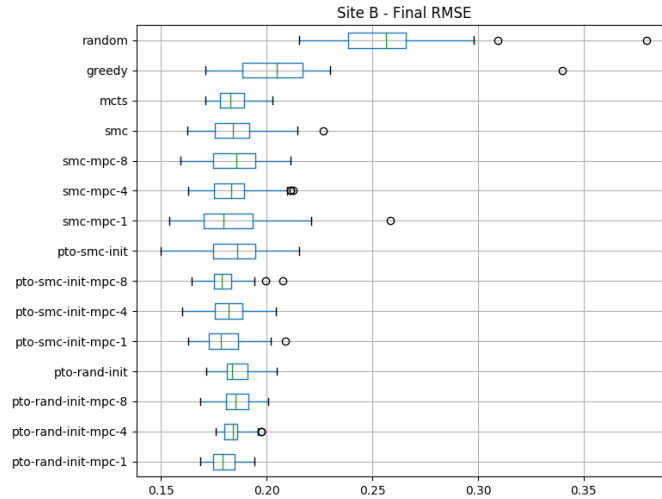


Figure 31: The statistical variation of final RMSE between runs is shown here. Once again, trends of the entropy performance are matched in the RMSE performance, particularly when comparing planning horizon within a planner type. Greedy notably has a much smaller IQR for RMSE compared to entropy.




Figure 32: 2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final RMSE performance between planners with 95% confidence. The entropy results are reflected here with notable differences. The best performing planner, PTO-SMC-INIT-MPC-1, is no longer statistically different than its competition, namely MCTS. This is more a function of the spectroscopic mapping model as the planners directly optimize over the entropy signal.

A.3 Site C




Figure 33: The smoothness of ergodic planner paths is shown well here. The MCTS path by comparison has many sharp turns and crosses back over itself. This example shows the Greedy planner performing well but the sharp back and forth about half way into the path shows the sensitivity of the planner to noise.




Figure 34: This example highlights the difference between PTO-SMC-INIT and PTO-RAND-INIT. Each planner converges to a different minima despite the same initial position however both are very effective at reducing entropy.




Figure 35: Average of all runs reported, lower entropy is better performance. Random performs the worst and performance improves as the planners become less myopic. Site C shows SMC not performing as closely to the best planners compared to the other sites.




Figure 36: Average of all runs reported, lower RMSE is better performance. Note that signal trends are highly correlated between entropy and RMSE, therefore our entropy calculation is a reliable information metric for use in informative path planning.




Figure 37: The statistical variation of final entropy between runs is shown here. There is a clear trend in median final entropy reducing as planning horizon reduces.




Figure 38: 2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence. PTO-SMC-INIT-MPC-1 performs best and is not statistically different than high performing planners MCTS, PTO-SMC-INIT-MPC-1, and PTO-SMC-INIT-MPC-4.




Figure 39: The statistical variation of final RMSE between runs is shown here. Trends of the entropy performance are matched in the RMSE performance.




Figure 40: 2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final RMSE performance between planners with 95% confidence. The entropy performance again is reflected. The statistical differences are similar only variants that differ by planning horizon are no longer statistically different.

A.4 Site D

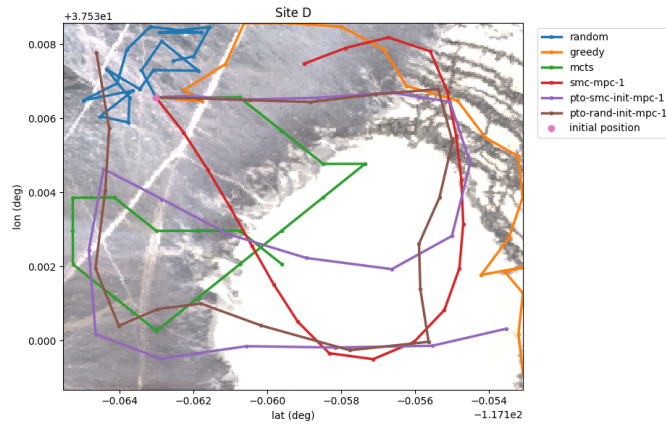


Figure 41: Site D is highly uniform in spectral diversity. The large white location is homogeneous in mineral type and therefore spectra signature. Note that planners effectively navigate the bi-model distribution. This site was chosen to investigate performance on regions with more uniform belief state, i.e. a less informative prior.

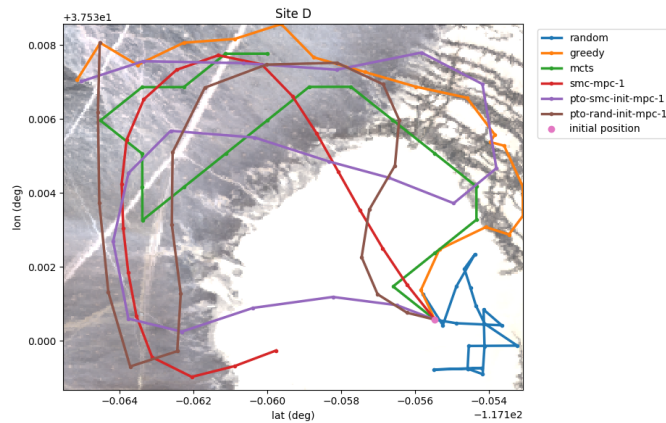


Figure 42: This example shows the planners spending more time in the grey locations with more information. PTO and SMC appear to more favorably balance exploration and exploitation but not aggressively avoiding the white location like MCTS.




Figure 43: Average of all runs reported, lower entropy is better performance. The performance difference between all planners is more pronounced in site D due to the more uniform information distribution. Note that all ergodic planners outperform other baselines indicating they do favorably balance exploration and exploitation.




Figure 44: Average of all runs reported, lower RMSE is better performance. Note that signal trends are highly correlated between entropy and RMSE, therefore our entropy calculation is a reliable information metric for use in informative path planning.

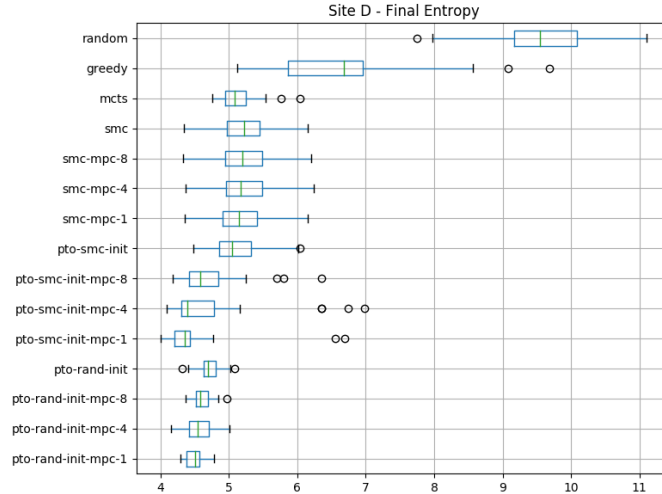


Figure 45: The statistical variation of final entropy between runs is shown here. Greedy has a large IQR given the myopic nature of the planner and the uniformity of the site. Note that the PTO-SMC-INIT variants contain more outliers than usual due to local minima in the uniform site.




Figure 46: 2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence. Nearly all of the PTO variants are statistically different from other baselines and outperform them including MCTS.




Figure 47: The statistical variation of final RMSE between runs is shown here. Once again, trends of the entropy performance are matched in the RMSE performance.




Figure 48: 2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final RMSE performance between planners with 95% confidence. The entropy and RMSE statistical difference is nearly identical with only PTO-SMC-INIT-MPC-8 no longer different than SMC-MPC-1 and PTO-SMC-INIT.

A.5 Site E

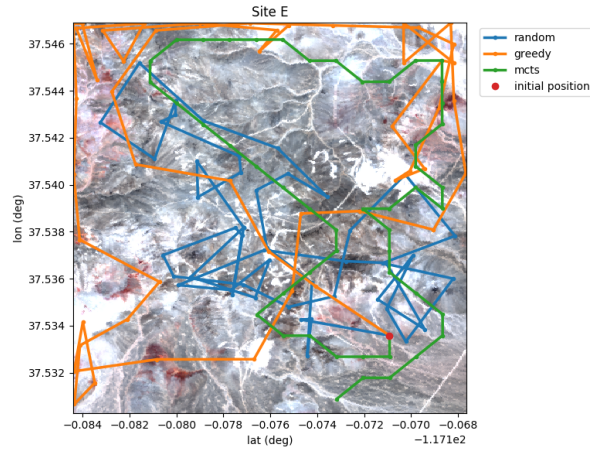


Figure 49: Only the baseline planner paths are shown here due to the long paths. Greedy is distracted by noise as is expected.

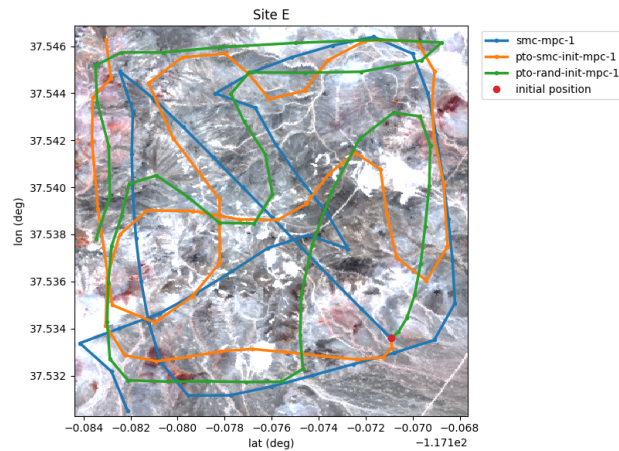


Figure 50: The ergodic planner paths are shown here. Their paths are notably smoother than the baseline planner paths. There is path crossover in the PTO-SMC-INIT path but the curves remain smooth. The SMC has sharp turns which is explained by significant changes to the entropy map after the previous sample.

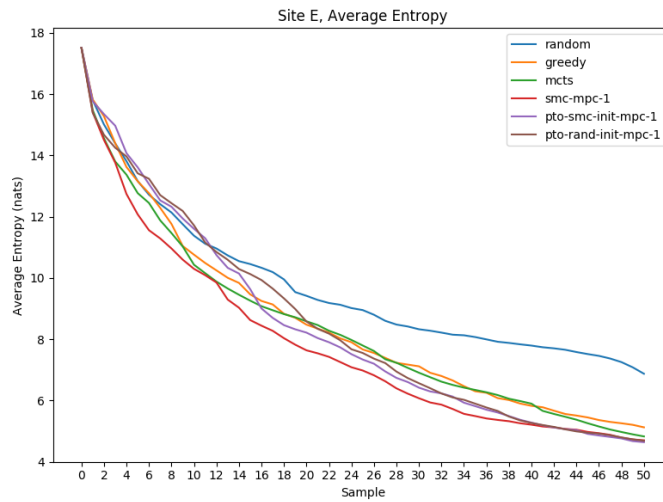


Figure 51: Average of all runs reported, lower entropy is better performance. Random performs the worst and performance improves as the planners become less myopic. Of significance here is SMC’s strong performance, outperforming MCTS throughout the entire average of the trajectories.




Figure 52: Average of all runs reported, lower RMSE is better performance. Note that signal trends are highly correlated between entropy and RMSE, therefore our entropy calculation is a reliable information metric for use in informative path planning.




Figure 53: The statistical variation of final entropy between runs is shown here. Interestingly PTO-SMC-INIT-MPC-8, PTO-RAND-INIT-MPC-8, and SMC-MPC-8 have large IQR or outliers while the other planning horizon variants do not. This suggests there are local minima modes in the entropy mode that are highlighted by updates after 8 steps but not finer or courser updates.




Figure 54: 2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final entropy performance between planners with 95% confidence. The majority of planner variants are statistically different from one another. Notably, the best performing planner PTO-SMC-INIT-MPC-1 is statistically different from all others except SMC-MPC-4 and PTO-RAND-INIT-MPC-1.




Figure 55: The statistical variation of final RMSE between runs is shown here. Once again, trends of the entropy performance are mirrored in the RMSE performance.




Figure 56: 2-sided dependent t-test for paired samples, cells that are dark indicate a statistically significant difference in final RMSE performance between planners with 95% confidence. Key differences between entropy performance and RMSE performance include the high performance of SMC variants. They are no longer statistically different from the highest performing planner, PTO-SMC-INIT-MPC-1.

References

- [1] Ian Abraham and Todd D. Murphey. “Decentralized Ergodic Control: Distribution-Driven Sensing and Exploration for Multiagent Systems”. In: *IEEE Robotics and Automation Letters* 3.4 (Oct. 2018), pp. 2987–2994. ISSN: 2377-3774. DOI: 10.1109/lra.2018.2849588. URL: <http://dx.doi.org/10.1109/LRA.2018.2849588>.
- [2] L. Armijo. “Minimization of functions having lipschitz continuous first partial derivatives”. In: *Pacific Journal of Mathematics*, vol. 16, no. 1, pp. 1-3, 1966.
- [3] E. Ayvali, H. Salman, and H. Choset. “Ergodic coverage in constrained environments using stochastic trajectory optimization”. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2017, pp. 5204–5210.
- [4] Elif Ayvali and Howie Choset. “Multi-Agent Ergodic Coverage with Obstacle Avoidance”. In: June 2017.
- [5] J. Binney, A. Krause, and G. S. Sukhatme. “Informative path planning for an autonomous underwater vehicle”. In: *2010 IEEE International Conference on Robotics and Automation*. 2010, pp. 4791–4796.
- [6] J. Binney and G. S. Sukhatme. “Branch and bound for informative path planning”. In: *2012 IEEE International Conference on Robotics and Automation*. 2012, pp. 2147–2154.
- [7] A. Candela et al. “Planetary robotic exploration driven by science hypotheses for geologic mapping”. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2017, pp. 3811–3818.

- [8] A. Candela et al. “Planetary Rover Exploration Combining Remote and In Situ Measurements for Active Spectroscopic Mapping”. In: *Proceedings of (ICRA) International Conference on Robotics and Automation*. May 2020.
- [9] Alberto Candela, David R. Thompson, and David Wettergreen. “Automatic Experimental Design Using Deep Generative Models of Orbital Data”. In: *Proceedings of The International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*. June 2018.
- [10] Nannan Cao, Kian Hsiang Low, and John M. Dolan. “Multi-Robot Informative Path Planning for Active Sensing of Environmental Phenomena: A Tale of Two Algorithms”. In: *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems*. AAMAS '13. St. Paul, MN, USA: International Foundation for Autonomous Agents and Multiagent Systems, 2013, pp. 7–14. ISBN: 9781450319935.
- [11] Doo-Hyun Cho et al. “Informative Path Planning and Mapping with Multiple UAVs in Wind Fields”. In: *Springer Proceedings in Advanced Robotics* (2018), pp. 269–283. ISSN: 2511-1264. DOI: 10.1007/978-3-319-73008-0_19. URL: http://dx.doi.org/10.1007/978-3-319-73008-0_19.
- [12] RN Clark. “Imaging spectroscopy: Earth and planetary remote sensing with the USGS Tetracoder and expert systems”. In: *J. Geophys. Res.* 108 (2003), pp. 1–44.
- [13] Roger N Clark et al. “Spectroscopy of rocks and minerals, and principles of spectroscopy”. In: *Manual of remote sensing* 3.3-58 (1999), pp. 2–2.
- [14] B. Cox. *Updated avionics electronic trends*. Tech. rep. 033. OPFMS, 2007.
- [15] N. A. Cressie. *Statistics for Spatial Data*. An optional note. John Wiley & Sons, Inc., Sept. 1993. ISBN: 9780471002550.

- [16] L. Dressel and M. J. Kochenderfer. “Tutorial on the generation of ergodic trajectories with projection-based gradient descent”. In: *IET Cyber-Physical Systems: Theory Applications* 4.2 (2019), pp. 89–100.
- [17] Greydon Foil. “Efficiently Sampling from Underlying Physical Models”. PhD thesis. Pittsburgh, PA: Carnegie Mellon University, Oct. 2016.
- [18] H. Fujisada et al. “Design and preflight performance of ASTER instrument protoflight model”. In: *IEEE Transactions on Geoscience and Remote Sensing* 36.4 (1998), pp. 1152–1160.
- [19] Bo-Cai Gao, Alexander FH Goetz, and JA Zamudio. “Removing atmospheric effects from AVIRIS data for surface reflectance retrievals”. In: (1991).
- [20] Emilio Garcia-Fidalgo and Alberto Ortiz. “Vision-based topological mapping and localization methods: A survey”. In: *Robotics and Autonomous Systems* 64 (2015), pp. 1–20.
- [21] S. Gautam et al. “Science-aware exploration using entropy-based planning”. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2017, pp. 3819–3825.
- [22] Robert O Green et al. “Imaging Spectroscopy and the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS)”. In: *Remote Sensing of Environment* 65.3 (1998), pp. 227–248. ISSN: 0034-4257. DOI: [https://doi.org/10.1016/S0034-4257\(98\)00064-9](https://doi.org/10.1016/S0034-4257(98)00064-9). URL: <http://www.sciencedirect.com/science/article/pii/S0034425798000649>.
- [23] L. Hamlin et al. “Imaging spectrometer science measurements for Terrestrial Ecology: AVIRIS and new developments”. In: (2011), pp. 1–7.
- [24] Dianyuan Han. “Comparison of commonly used image interpolation methods”. In: (2013).

- [25] John Hauser. “A PROJECTION OPERATOR APPROACH TO THE OPTIMIZATION OF TRAJECTORY FUNCTIONALS”. In: *IFAC Proceedings Volumes* 35.1 (2002). 15th IFAC World Congress, pp. 377–382. ISSN: 1474-6670. DOI: <https://doi.org/10.3182/20020721-6-ES-1901.00312>. URL: <http://www.sciencedirect.com/science/article/pii/S1474667015387334>.
- [26] Daniel C Heinz et al. “Fully constrained least squares linear spectral mixture analysis method for material quantification in hyperspectral imagery”. In: *IEEE transactions on geoscience and remote sensing* 39.3 (2001), pp. 529–545.
- [27] Geoffrey A. Hollinger and Gaurav S. Sukhatme. “Sampling-based robotic information gathering algorithms”. In: *The International Journal of Robotics Research* 33.9 (2014), pp. 1271–1287. DOI: 10.1177/0278364914533443.
- [28] I. I. Hussein and D. M. Stipanovic. “Effective Coverage Control for Mobile Sensor Networks With Guaranteed Collision Avoidance”. In: *IEEE Transactions on Control Systems Technology* 15.4 (2007), pp. 642–657.
- [29] Alonzo Kelly. *Mobile robotics: mathematics, models, and methods*. Cambridge University Press, 2013.
- [30] Diederik P Kingma and Max Welling. *Auto-Encoding Variational Bayes*. 2013. arXiv: 1312.6114 [stat.ML].
- [31] D.E. Kirk, D.N. Kirk, and D.L. Kreider. *Optimal Control Theory: An Introduction*. Networks series. Prentice-Hall, 1970.
- [32] S. Kodgule, A. Candela, and D. Wettergreen. “Non-myopic Planetary Exploration Combining In Situ and Remote Measurements”. In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2019, pp. 536–543.

- [33] Andreas Krause, Ajit Singh, and Carlos Guestrin. “Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies”. In: *J. Mach. Learn. Res.* 9 (June 2008), pp. 235–284. ISSN: 1532-4435.
- [34] Fred A Kruse and Sandra L Perry. “Regional mineral mapping by extending hyperspectral signatures using multispectral data”. In: (2007), pp. 1–14.
- [35] Zhan Wei Lim, David Hsu, and Wee Sun Lee. “Adaptive informative path planning in metric spaces”. In: *The International Journal of Robotics Research* 35.5 (2016), pp. 585–598. DOI: 10.1177/0278364915596378.
- [36] R. Marchant and F. Ramos. “Bayesian Optimisation for informative continuous path planning”. In: *2014 IEEE International Conference on Robotics and Automation (ICRA)*. 2014, pp. 6136–6143.
- [37] George Mathew and Igor Mezić. “Metrics for ergodicity and design of ergodic dynamics for multi-agent systems”. In: *Physica D: Nonlinear Phenomena* 240.4 (2011), pp. 432–442. ISSN: 0167-2789. DOI: <https://doi.org/10.1016/j.physd.2010.10.010>. URL: <http://www.sciencedirect.com/science/article/pii/S016727891000285X>.
- [38] L. M. Miller and T. D. Murphey. “Trajectory optimization for continuous ergodic exploration”. In: *2013 American Control Conference*. 2013, pp. 4196–4201.
- [39] L. M. Miller et al. “Ergodic Exploration of Distributed Information”. In: *IEEE Transactions on Robotics* 32.1 (2016), pp. 36–52.
- [40] Lauren Miller. “Optimal ergodic control for active search and information acquisition”. PhD thesis. Northwestern University, 2015.

- [41] P. Morere, R. Marchant, and F. Ramos. “Sequential Bayesian optimization as a POMDP for environment monitoring with UAVs”. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. 2017, pp. 6381–6388.
- [42] Scott Murchie et al. “Compact reconnaissance imaging spectrometer for Mars (CRISM) on Mars reconnaissance orbiter (MRO)”. In: *Journal of Geophysical Research: Planets* 112.E5 (2007).
- [43] John F Mustard, Carle M Pieters, and Stephen F Pratt. “Deconvolution of spectra for intimate mixtures”. In: 17 (1986), pp. 593–594.
- [44] Carlo Orrieri, Gianmario Tessitore, and Petr Veverka. “Ergodic Maximum Principle for Stochastic Systems”. In: *Applied Mathematics & Optimization* 79.3 (Sept. 2017), pp. 567–591. ISSN: 1432-0606. DOI: 10.1007/s00245-017-9448-7. URL: <http://dx.doi.org/10.1007/S00245-017-9448-7>.
- [45] K. Petersen. *Ergodic Theory*. Cambridge University Press, 1984.
- [46] Shujun Qiao et al. “A review of laser-induced breakdown spectroscopy for analysis of geological materials”. In: *Applied Spectroscopy Reviews* 50.1 (2015), pp. 1–26.
- [47] CE. Rasmussen and CKI. Williams. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. Cambridge, MA, USA: MIT Press, Jan. 2006, p. 248.
- [48] D. A. Roberts et al. “Temporal and spatial relationships between topography, atmospheric water vapor, liquid water and vegetation endmember fractions determined using AVIRIS”. In: 4 (1994), 2366–2368 vol.4.

- [49] Juanita C Sandidge and Ronald J Holyer. “Coastal bathymetry from hyperspectral observations of water radiance”. In: *Remote Sensing of Environment* 65.3 (1998), pp. 341–352.
- [50] Michael C Shewry and Henry P Wynn. “Maximum entropy sampling”. In: *Journal of applied statistics* 14.2 (1987), pp. 165–170.
- [51] Y. Silverman et al. “Optimal planning for information acquisition”. In: *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2013, pp. 5974–5980.
- [52] Geoffrey M Smith and Edward J Milton. “The use of the empirical line method to calibrate remotely sensed data to reflectance”. In: *International Journal of remote sensing* 20.13 (1999), pp. 2653–2662.
- [53] Gregg Swayze et al. “Ground-truthing AVIRIS mineral mapping at Cuprite, Nevada”. In: (1992).
- [54] Gregg A Swayze et al. “Mapping advanced argillic alteration at Cuprite, Nevada, using imaging spectroscopy”. In: *Economic Geology* 109.5 (2014), pp. 1179–1221.
- [55] D. R. Thompson et al. “On optimal estimation theory for atmospheric correction in vswir imaging spectroscopy”. In: (2017), pp. 3063–3065.
- [56] David R Thompson et al. “Spatial spectroscopic models for remote exploration”. In: *Astrobiology* 18.7 (2018), pp. 934–954.
- [57] David R. Thompson. “Intelligent Mapping for Autonomous Robotic Survey”. PhD thesis. Pittsburgh, PA: Carnegie Mellon University, 2008.
- [58] David Ray Thompson et al. “Spatio-Spectral Exploration Combining In Situ and Remote Measurements.” In: *AAAI*. 2015, pp. 3679–3685.
- [59] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. Cambridge, Mass.: MIT Press, 2005. ISBN: 0262201623.

- [60] C. Veitch, D. Render, and A. Aravind. “Ergodic Flocking”. In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2019, pp. 6957–6962.
- [61] C. Wong et al. “Adaptive and intelligent navigation of autonomous planetary rovers — A survey”. In: *2017 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*. 2017, pp. 237–244.
- [62] Himanshi Yadav, Alberto Candela, and David Wettergreen. “A Study of Unsupervised Classification Techniques for Hyperspectral Datasets”. In: (2019), pp. 2993–2996.