2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Resource Mapping with a Mobile Exploration Robot using Spectral Mixture Ergodic Search

Margaret Hansen[1*], Ananya Rao[1], Abigail Breitfeld[1], David Wettergreen[1]

***Abstract*— Resource mapping and prospecting has become the focus of a number of proposed planetary exploration missions, particularly to locate water ice at the lunar south pole. Mobile robots, which are employed for exploration tasks in environments that are inaccessible to humans, collect the information in such missions. In these scenarios, intelligent and adaptive trajectory planning algorithms increase the accuracy of the resulting resource map, along with the efficiency with which information is gathered.**

In this work, we use ergodic search to generate a mobile robot trajectory that balances exploration and exploitation, while simultaneously mapping the spatial distribution of a resource by using Gaussian process regression with a spectral mixture kernel. The spatial correlation structure learned via Gaussian process regression informs the ergodic search about regions of high information, as well as the frequency components that appear in the map distribution. We call this method spectral mixture ergodic search (SM-ES) and demonstrate how it learns a map and updates the trajectory accordingly on three datasets: synthetic maps, an ice favorability index map for the lunar south polar region, and real mineral data from Cuprite, Nevada.

I. Introduction

Resource mapping is a crucial step towards using off-world resources for future space missions. Understanding the distribution of cold-trapped ice on the moon is listed as a high priority for the Artemis program's science goals, and NASA's Endurance seeks to characterize the geologic diversity across the South Pole–Aitken basin [1,2]. Other nations have also developed missions to the lunar south pole, such as China's Yutu-2 rover and India's Pragyan [3,4]. Mobile exploration robots are used to investigate extreme environments (space, remote deserts, undersea, etc.) that are too dangerous or inaccessible to humans—among other things, these robots can be tasked with geologic and mineralogic surveying, search and rescue, and radiological monitoring [5–8]. As a result, recent focus among the planetary science community has been on using robots for resource mapping and prospecting, particularly for investigating the distribution of water ice in the south polar region of the moon through programs such as the European Space Agency's PROSPECT payload, NASA's PRIME-1 experiment, NASA's Volatiles Investigating Polar Exploration Rover (VIPER) rover, and commercial ventures such as Astrobotic's MoonRanger [9–12]. Resource mapping robots can also be deployed on other planetary bodies with potential resources, such as on Mars or on asteroids [13,14].

In this work, Gaussian process regression is paired with ergodic search to iteratively update a map and re-plan based on the map updates in a manner that enables the adaptation of the search hyperparameters. This process is demonstrated in Figure 1, which displays the iterative operational cycle in which frequencies learned from the map are used in the trajectory optimization step, followed by map and frequency updates based on new observations. Gaussian process regression with a spectral mixture kernel is used for mapping due to its ability to incorporate spatial correlations, enable extrapolation to unobserved regions, and estimate frequencies that best describe a map's spatial correlation structure. In addition, the predictive uncertainty from the GP model represents the importance of sampling in a given region and

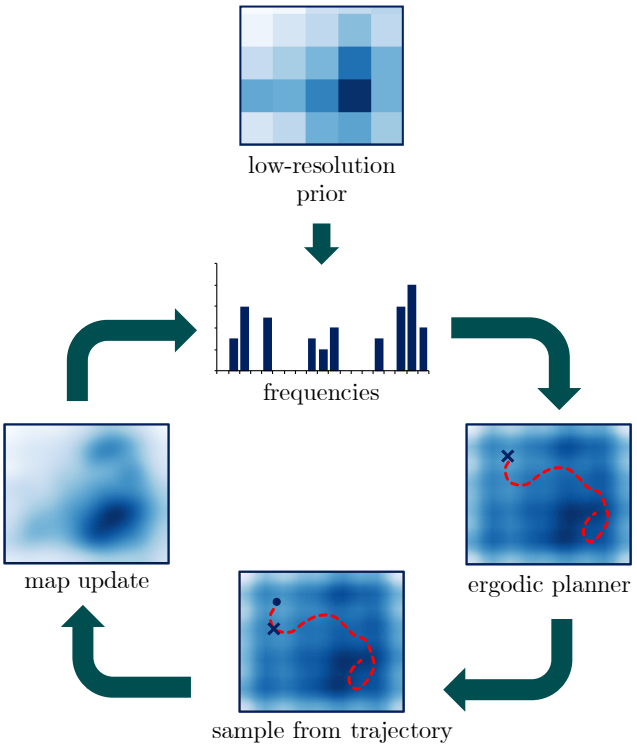

Fig. 1: SM-ES iteratively learns the map and updates the trajectory based on the frequencies found to best describe the map's spatial correlation structure. The ground truth map in this example is a multimodal Gaussian, where blue indicates higher information and white indicates lower information. An initial ergodic trajectory (red) is planned based on predictions generated from a low-resolution prior. After a sample is taken along the trajectory, the map is updated, and the frequencies from this new map are used to plan a new trajectory.

[1] Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA.
[*] Corresponding author. margareh@cs.cmu.edu.

is taken as the information distribution for use with ergodic search. Ergodic search is used to plan paths that balance exploration of regions with low information and exploitation of areas with high information via frequency space analysis of the map along with proposed trajectories [15]. Spectral mixture ergodic search (SM-ES) links the frequencies used in ergodic search to those found by the spectral mixture kernel in order to enable iterative adaptation of the search algorithm to the learned features of the map, leading to faster and more adaptable planning. SM-ES has been tested in simulation as well as with two real-world datasets: an ice favorability index map for the lunar south polar region [16] and mineral maps from Cuprite, Nevada [17].

II. Background

Both Gaussian process-based mapping and ergodic search have a history of use for information gathering and other tasks similar to resource mapping. Gaussian process regression has long been used to model quantities with spatial and/or temporal correlations, such as temperature or climate-ice sheet interactions [18,19]. Recent developments in robotics have relied on this method for modeling information distributions due to its ability to provide the out-of-sample predictions and uncertainties required for informative path planning [5,8]. Ergodic search, a search method that leverages a known information distribution to plan trajectories, is capable of balancing exploration of the search domain and exploitation of the available prior information to improve performance for tasks such as information gathering and disaster relief [20,21]. However, the two techniques have mainly developed separately from one another.

A. Resource Mapping

Resource *prospecting* is a branch of geologic science that involves searching for minerals or ores that can lead to the location of mineral deposits [22]. Within resource prospecting, which can include *in situ* sampling of subsurface mineral content with a drill or other equipment, resource *mapping* focuses on predicting the distribution of a resource before such samples are taken to map out high-probability regions for exploration [23]. The goal of resource mapping is to explore different regions and take samples to identify or confirm locations of interest for future resource prospecting. As such, it can be formulated as an information gathering task in which the gathered information becomes a map of resource locations and/or quantities [24]. Many approaches use remote sensing or expert information to generate resource maps [25,26]; such maps are limited to low-resolution observations and do not incorporate spatial correlations.

Gaussian process regression, termed "kriging" in the geostatistics community, is a commonly-used alternative that can model spatial correlations and extrapolate outside of observed regions. This approach models the quantity of interest as a function of 2D location with additive noise and a Gaussian process (GP) prior parameterized by a mean function and a kernel function [27]. In the context of resource mapping, mineral prospectivity maps have relied on kriging to interpolate geochemical or geophysical properties [23]. Robotic mapping methods have frequently used GPs to model information distributions for informative path planning due to their ability to provide uncertainty estimates [8,17].

B. Ergodic Search

Ergodic search is a method of trajectory optimization that minimizes the *ergodic metric*: the difference between the information observed along a proposed trajectory and the information map being explored [15]. Ergodic trajectories spend time in areas of the map proportional to the expected amount of information there, thereby balancing exploitation of regions with high information with exploration of regions with low information [15], and have been shown to be effective for information gathering in a variety of settings [28–30]. Ergodic search has previously been paired with kernel-based methods to extend it to arbitrary search domains [31] or to increase computational efficiency [32].

While ergodic search is optimal under certain conditions and assumptions, such as submodularity (in which repeated measurements at a given location provide successively less information), the basic method does not consider spatial correlations present in the underlying information distribution [33]. Moreover, the prior literature relies on static, predetermined hyperparameters in optimization, even for dynamic information distributions [34]. Works that have incorporated kernel methods do not focus on mapping information as a region is explored. The current work addresses this gap in the literature by adapting the set of frequencies to the search domain as an information distribution is learned.

To our knowledge, no prior work attempts to mathematically pair robotic mapping with ergodic search such that the mapping procedure learns parameters directly relevant to the computation of the metric used in trajectory optimization. By connecting the two in this way, the parameters of ergodic search are made more adaptive to the spatial distribution being learned. This extension also enables learning different frequencies in different spatial dimensions, allowing the trajectory optimization to use frequencies relevant to a given dimension based on the spatial correlation structure of the information distribution.

III. Methodology

Spectral mixture ergodic search (SM-ES) combines Gaussian process-based map estimation and ergodic search by iteratively updating the map and re-planning until the trajectory length has been exhausted, following the steps shown in Figure 1. The map is initialized using low-resolution prior observations, with the learned frequencies and weights used to inform ergodic search. The planning step then relies on these parameters when optimizing the trajectory such that the full trajectory minimizes the ergodic metric. After the robot moves and takes an observation, the map is updated with the new data, and the trajectory is re-optimized. Iteration occurs until the robot has taken a pre-specified number of

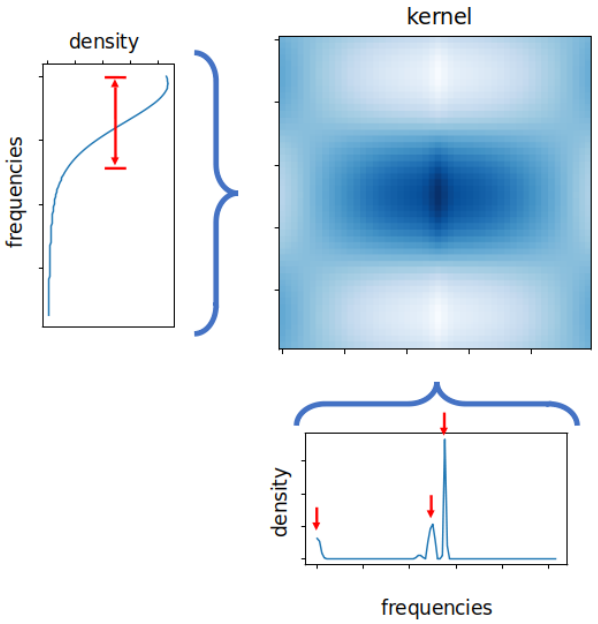

Fig. 2: SM-ES relies on the spectral mixture kernel, which learns important frequencies for each input dimension for modeling the spatial correlation structure of the data. In this example, a spectral mixture kernel learned from one of the synthetic data maps is decomposed into a Gaussian mixture of frequencies in each spatial dimension. The kernel has both low and high frequency components in the horizontal dimension (red arrows in the bottom plot), indicating correlations specific to certain distances. In the vertical dimension (red interval in the left plot), this kernel exhibits only low, overlapping frequency components—albeit with higher variances—corresponding to a more steady, gradual decrease in correlation with greater distance.

steps. Since we are using ergodic search in an information gathering context, the distribution used for planning is the uncertainty derived from the resource map as quantified by the variance associated with the predicted map. For our experiments, the map is initialized using a prior with low spatial resolution to simulate orbital satellite data that would likely be available for applications such as prospecting for water on the moon or minerals on Mars.

A. *Learning the Map*

For this work, a resource map displays the probability p that a resource is present at a given location $\mathbf{x} \in \mathbb{R}^2$. To incorporate spatial correlations within the resource map, Gaussian process regression is used to model p as a function of location with a Gaussian process prior with mean function $m(\mathbf{x})$ and kernel function $k(\mathbf{x}, \mathbf{x}')$: $p(\mathbf{x}) \sim GP(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$. The resource map then becomes the predicted probability $\hat{p}$ at grid points across the search domain based on the Gaussian process model.

In learning the map, noisy observations of the resource probability $y_i = p(\mathbf{x}_i) + \varepsilon$ are made at locations $\mathbf{x}_i$ for $i \in \{1, ..., N\}$. Given the training dataset $(X, \mathbf{y}) = \{(\mathbf{x}_i, y_i)\}_i$ and a noise model $\varepsilon \sim N(0, \sigma_n^2)$, the Gaussian process regression model can predict the expected value and variance of p at any query location $\mathbf{x}^*$ according to the well-known analytic equations shown in Equations 1a and 1b, with $\mathbf{k} = k(\mathbf{x}^*, X)$ and $K = k(X, X)$ used to simplify the notation [27]:

$$\hat{p}(\mathbf{x}^*) = m(\mathbf{x}^*) + \mathbf{k}^T (K + \sigma_n^2 I)^{-1} (\mathbf{y} - m(X)) \quad (1a)$$

$$\mathrm{Var}(\hat{p}) = k(\mathbf{x}^*, \mathbf{x}^*) - \mathbf{k}^T (K + \sigma_n^2 I)^{-1} \mathbf{k} \quad (1b)$$

As the robot moves through a trajectory, more points are added to the training dataset and the resource map is updated to reflect the new information. The uncertainty associated with the new resource map is computed using Equation 1b and provided to the planner to re-plan the remainder of the trajectory based on the new observation.

Two key design decisions for Gaussian process regression are the choice of mean function $m(\mathbf{x})$ and kernel function, $k(\mathbf{x}, \mathbf{x}')$, along with any associated hyperparameters. The mean function is frequently set to zero or a constant; we use the latter in this work and learn the constant during the optimization step. The kernel function k used in this work is the spectral mixture kernel, which is the inverse Fourier transform of a Gaussian mixture model (GMM) and has the closed form shown in Equation 2 [35,36]:

$$k(\mathbf{d}) = \sum_{q=1}^{Q} w_q \cos\left(2\pi \mathbf{d}^T \boldsymbol{\mu}_q\right) \prod_{m=1}^{M} \exp\left(-2\pi^2 \mathbf{d}_m^2 \boldsymbol{\nu}_q^{(m)}\right) \quad (2)$$

In this equation, $\mathbf{d} = \mathbf{x} - \mathbf{x}'$ is the difference between the two locations $\mathbf{x}$ and $\mathbf{x}'$, Q is the number of mixture components in the GMM, and M represents the dimensionality of the input $\mathbf{d}$ ($M = 2$ for the 2-dimensional maps used here). The hyperparameters of this kernel are w_q, $\boldsymbol{\mu}_q$, and $\boldsymbol{\nu}_q$, which represent the weight, mean vector, and variance vector,[1] respectively, of the qth component of the Gaussian mixture. Since the GMM is defined in the frequency domain, the mixture component means $\boldsymbol{\mu}_q$ represent frequencies estimated from the data, as shown in Figure 2.

The spectral mixture kernel was chosen for a few reasons. First, spectral mixture kernels are quite expressive and can approximate *any* stationary kernel function (i.e. a kernel that depends only on the distance between inputs $\mathbf{d}$, which includes common kernels such as the radial basis function and Matérn kernels) to arbitrary precision, depending on the number of GMM components [27,35]. As a result, the spectral mixture kernel is more expressive and flexible than standard kernels, hence its use for all of our tests. Additionally, the mean frequency vectors $\boldsymbol{\mu}_q$ and component weights w_q learned by the kernel are directly relevant to ergodic search and are used by SM-ES in the computation of the ergodic metric.

In our implementation, the mean function and kernel hyperparameters are optimized by maximizing the likelihood of

[1]The covariance matrices for the Gaussian mixture model components are assumed to be diagonal with $\boldsymbol{\nu}_q$ representing the diagonal elements [35].

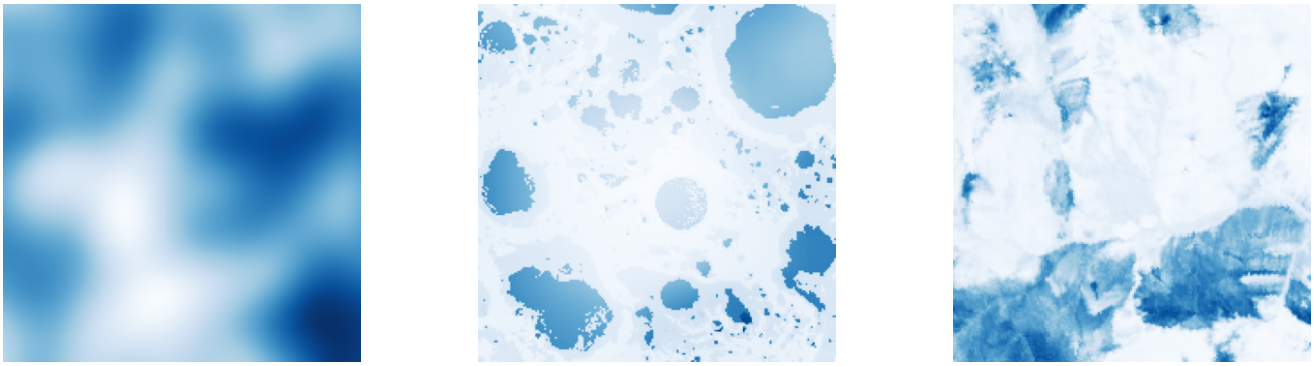

(a) 20-Frequency Synthetic Map (b) Ice Favorability Index (c) Opal Mineral Map

Fig. 3: Maps corresponding to each of the datasets used for experiments. The synthetic map (a) and Cuprite (Opal) map (c) are example maps from datasets with 5 and 6 maps, respectively, all of which were used for experiments. The ice favorability index (b) is the only map for that dataset.

the data using the Adam optimizer [37] and GPyTorch [38], which contains an implementation of the spectral mixture kernel. The frequencies are capped to $1/2r$ based on the resolution of the predicted map r to avoid aliasing [39]. Exact inference is used for all maps; future work will address incorporating approximate inference to avoid the scalability issues associated with exact GP inference. Map updates are performed after each step in the trajectory by adding the new observation to the dataset and re-optimizing, starting with the previous hyperparameter values. Based on initial testing with the mapping method, the number of components used by the spectral mixture kernel for map prediction is set to 10 for all experiments based on initial testing with this data. Fewer than 10 results in maps with substantially higher map error, while increasing the number of components beyond 10 results in a minimal error decrease but increase in training and inference time.

B. Spectral Mixture Ergodic Search

Spectral mixture ergodic search (SM-ES) uses the standard formulation of the ergodic metric but with frequencies and weights that are learned from the map, as shown in Equation 3:

$$\mathcal{E}(\mathbf{x}(t)) = \sum_{q=0}^{Q} \Lambda_q w_q ||c_q(\mathbf{x}(t)) - \phi_q||^2 \tag{3}$$

The ergodic metric represents a weighted average of the squared difference between the Fourier coefficients of the map (ϕ_q) and those of the distribution representing the time-averaged trajectory ($c_q(\mathbf{x}(t))$)—in simple terms, the difference between the information in the map and what information the trajectory expects to see [20]. The typical approach to ergodic search uses a standard set of Q frequency vectors across both spatial domains when computing the Fourier coefficients [15,20]. By contrast, SM-ES relies on frequencies learned from the map $\boldsymbol{\mu}_q$ in the computation of c_q and ϕ_q, along with the corresponding weights w_q, which are used alongside the weights typically included in ergodic search, $\Lambda_q = (1+||\boldsymbol{\mu}_q||^2)^{-3/2}$ for 2-dimensional trajectories [21]. This combination of weights balances information from the map in the form of w_q with ergodic search's standard bias towards lower frequency information to promote exploration and avoid high-frequency noise. The trajectory optimization proceeds as a minimization of the ergodic metric subject to dynamics constraints of the robot and boundary conditions. This work assumes a forward-moving ground-based robot with differential drive and optimizes the linear and angular velocities. Soft constraints are imposed on the velocities with a penalty that scales with the velocity magnitude, and a boundary constraint that penalizes trajectories that leave the map region. PyTorch's Adam optimizer is used to learn the robot controls subject to these constraints [37].

Standard ergodic search requires definition of frequencies prior to learning map information; in practice, this is done using all combinations across the same range of frequencies in multiple dimensions. For example, a 2D ergodic search trajectory using 10 frequencies per dimension relies on 100 frequency pairs. SM-ES, on the other hand, learns frequency pairs from the map. When using 10 map components, only 10 pairs are learned, and these frequencies may be different depending on the information in each dimension. Additionally, these parameters change over time as the map is learned, meaning that the SM-ES algorithm itself is able to adapt to changes in the underlying map. Our version of the ergodic search algorithm is loosely based on the implementation of multi-objective ergodic search (MO-ES) [21], modified to use PyTorch as the optimization framework and to use frequency and weight inputs from the map. This ergodic search implementation is available at `https://github.com/PlanetaryRobotics/ergodic-search`.

IV. Experiments

Spectral mixture ergodic search has been tested on synthetic data and two real-world datasets: an ice favorability map for the lunar south polar region and mineral maps of a region near Cuprite, Nevada. Examples from each dataset

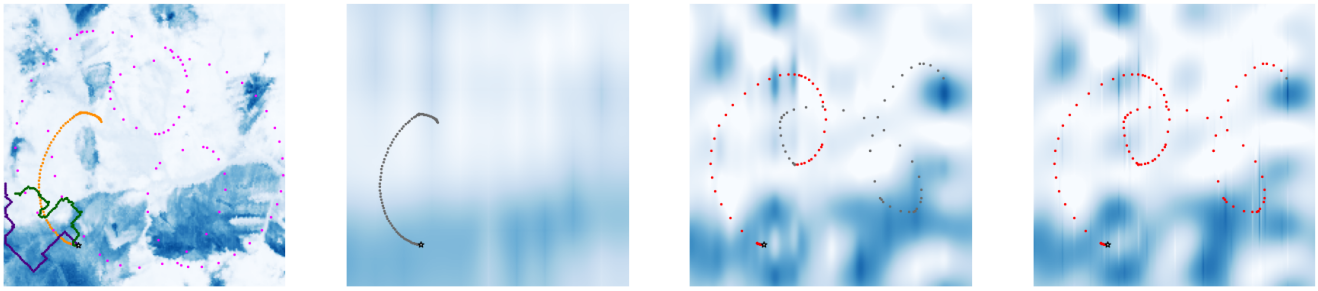

(a) (b) (c) (d)

Fig. 4: Example from the Cuprite experiments for predicting the presence of the mineral opal. (a) The ground truth map shows the probability that the mineral opal is present in a given location, along with the final trajectories generated by the Greedy planner (indigo), MCTS (green), Prior ES (orange), and Baseline ES (magenta). (b) SM-ES initially has some estimate of where the opal is from the low-resolution prior, but these regions lack detail. The initial SM-ES trajectory (black) matches the Prior ES trajectory. (c) After 50 steps, the predicted map looks more similar to the ground truth and the trajectory has shifted to explore additional regions. Red indicates the SM-ES trajectory taken prior to the current step, and black indicates the planned SM-ES trajectory. (d) The final predicted map has higher frequency features and more closely resembles the ground truth. The red trajectory is the completed SM-ES trajectory.

are shown in Figure 3. The synthetic data is used for proof-of-concept, while the ice favorability map and Cuprite experiments demonstrate applicability to real-world datasets for resource mapping scenarios.

A. Datasets

1) Synthetic Maps: The synthetic dataset consists of five resource probability maps that have been generated as mixtures of random noise with a range of specified frequencies, with the simplest map combining 5 frequencies and the most complicated combining 50. When generating the low-resolution prior and samples from these synthetic maps, Gaussian noise with a standard deviation of 0.02 percentage points was added to simulate sensor noise. The ground truth maps were created at a resolution of 0.05 m per pixel, with a resolution of 10 m per pixel used for the prior and 1 m per pixel used for the predicted map.

2) Ice Favorability Index: The ice favorability index (IFI) map is included to represent a scenario similar to proposed missions to the lunar south polar region to prospect for lunar volatiles [9,12]. This map consists of an index indicating favorable conditions for subsurface volatiles, computed based on information such as the thermal environment of the surroundings [16]. This map covers the lunar south polar region from 80° S to 90° S latitude at a resolution of roughly 591 meters per pixel; for our experiments, the roughly 100 km x 100 km portion on either side of the south pole was extracted to provide a region that was not prohibitively large for exploration. The low-resolution prior was provided at roughly 11 km per pixel, while the predicted map was modeled at the same resolution as the ground truth.

3) Cuprite Mineral Maps: The Cuprite data consists of six mineral maps processed from hyperspectral images taken by the Next Generation Airborne Visible Infrared Imaging Spectrometer (AVIRIS-NG) [17]. These maps represent the probability that a mineral is present at a given location for six different minerals and were generated from the underlying hyperspectral images using a logistic regression classifier trained on data from the USGS Spectral Library Version 7 [40,41]. A test region that covers roughly 0.5 sq km was chosen from the larger images. The hyperspectral images have a resolution of 3.9 m per pixel [17]; the low-resolution prior was chosen to be 75 m per pixel and the predicted map was generated at the full resolution of 3.9 m per pixel.

B. Comparison Methods

SM-ES is compared to five planning methods for all three datasets. The first, the "**Full Coverage**" planner, provides a trajectory with complete coverage over the map based on a standard step size starting at one corner and ending at the diagonally opposite one. This planner does not consider the information content of the map but demonstrates a simple but inefficient coverage method that will provide a lower bound for performance in terms of map error. One full coverage trajectory was completed per map with a step size designed to get near 100 sampling points.

The other four planners are all informative path planners since they plan over the information distribution:

- **Greedy**: One-step trajectory that samples the best neighboring grid cell in an 8-connected grid after each map update.
- **Monte Carlo Tree Search (MCTS)**: Randomized grid search based on an 8-connected grid with information gain considered across ten planned future actions. Only the first planned action is executed before re-planning occurs.
- **Prior Ergodic Search (Prior ES)**: Ergodic search performed on the initial low-resolution prior without trajectory re-planning.
- **Baseline Ergodic Search (Baseline ES)**: Ergodic search with the standard set of frequencies and weights with trajectory re-planning after each map update.

	Planning Time (s)	MAE (↓)	Avg Var (↓)	ΔMAE$\times 10^{-5}$ (↓)	ΔVar$\times 10^{-5}$ (↓)
Synthetic Data					
Full Coverage	6.08E-5	0.0169	0.0007	-106.0	-26.3
Greedy	**0.0140**	0.0756	0.0047	-44.7	-19.8
MCTS	0.0545	0.0698	0.0034	-50.6	-21.2
Prior ES	1.6609	0.0888	0.0042	-38.6	**-24.7**
Baseline ES	0.9615	**0.0473**	**0.0024**	**-73.3**	-22.2
SM-ES	0.6826	0.0639	0.0032	-56.0	-21.7
Ice Favorability Index Data					
Full Coverage	5.13E-5	0.0925	0.0083	-28.73	-8.55
Greedy	**0.015**	0.1128	0.1090	-7.11	-4.92
MCTS	0.066	0.1124	0.0107	-7.50	-5.10
Prior ES	1.768	0.1063	0.0088	-15.11	-6.24
Baseline ES	1.257	0.1017	0.0085	-18.29	-7.38
SM-ES	1.396	**0.0999**	**0.0063**	**-20.16**	**-9.88**
Cuprite Data					
Full Coverage	5.32E-5	0.1457	0.0089	-3.86	-3.21
Greedy	**0.0150**	0.1485	0.0085	-1.10	-3.73
MCTS	0.0659	0.1484	0.0080	-1.22	-4.31
Prior ES	2.1920	0.1477	0.0082	-2.72	**-5.41**
Baseline ES	1.1854	**0.1466**	0.0082	**-3.25**	-4.07
SM-ES	1.1152	0.1470	**0.0078**	-2.78	-4.23

TABLE I: Average performance metrics across all starting points and maps for each of the three datasets. Bold indicates the lowest value for that metric in that dataset, across all planners except the full coverage planner.

The greedy and MCTS planners are included as alternatives to ergodic search that are commonly found in the informative path planning literature [42]. The baseline ergodic search method is included to show how the typical set of frequencies and weights performs in comparison to SM-ES. The prior ergodic search method is similar but only incorporates the low-resolution prior information.

For all informative path planners, 50 random starting locations were selected for each dataset. Each of the search algorithms was run from these starting locations for 100 steps. The three ergodic search methods were given the same parameters to ensure comparability, with optimization ending after the ergodic metric dropped below 0.01 or 1000 iterations were reached.

C. Metrics

The primary metrics used to evaluate SM-ES are the mean absolute error (MAE) of the final predicted resource map and the average planning time per iteration. The former measures the accuracy of the resource map as compared to a ground truth map after exploration is complete, while the latter measures efficiency of the planner. Secondary metrics are also measured, consisting of the average variance of the final resource map and the average decrease in MAE and map variance over each trajectory step. The latter two metrics can provide a sense of the rate at which a planner gathers information when exploring, while the former provides a sense of the uncertainty remaining after exploration.

V. Results & Discussion

Our experiments demonstrate that SM-ES exhibits performance that is overall on par with other informative path planners, though baseline ergodic search tends to still produce better overall accuracy metrics. Averages for all planners and metrics are organized by dataset in Table I. Baseline ergodic search achieves the lowest map error (aside from the full coverage method used as a lower bound) for two out of three maps, with SM-ES producing the lowest map error for the ice favorability index—though this is not a statistically significant difference (see Table II below).

The ergodic planners are all slower than the non-ergodic planners (Greedy, MCTS, Full Coverage); however, the ergodic planners account for robot dynamics in ways that the latter do not, as can be seen in the example trajectories in Figure 4. In this example, Greedy and MCTS produce trajectories that step along the grid one pixel at a time and can only make 45° or 90° turns at any given point. The ergodic planners, on the other hand, are capable of producing gentler curves and end up exploring more of the space. The SM-ES trajectory also follows an interesting pattern: the first trajectory is identical to the Prior ES trajectory, as is expected, but as the robot explores, it progresses towards a trajectory similar to the Baseline ES trajectory.

The remaining metrics exhibit similar trends. SM-ES achieves the lowest final average map variance for two out of three datasets, while Baseline ES performs the best in terms of the average decrease in MAE. Prior ES exhibits the greatest average decrease in map variance for two out of three maps. However, the decreases in MAE and variance from one iteration to the next are quite small (on the order of 10^{-5} to 10^{-4}), so these differences are practically negligible.

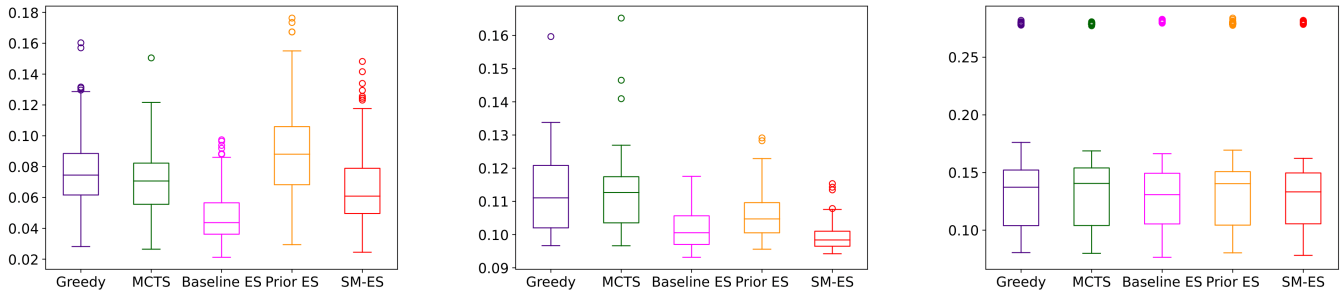

(a) Synthetic Maps (b) Ice Favorability Index (c) Cuprite Mineral Maps

Fig. 5: Box plots showing the final mean absolute error (MAE) for all starting points, for each of the three datasets. For the synthetic maps and the Cuprite mineral maps, results are averaged over all maps included in the respective datasets.

A. *Statistical Significance*

Wilcoxon signed rank tests, chosen to avoid assuming normality of the distributions [43], were performed to compare SM-ES to each of the other informative path planners on the primary metrics (planning time and final map error). The results of each of these paired tests is shown in Table II. SM-ES is statistically significantly faster than Prior ES on average for all three datasets, and statistically significantly faster than Baseline ES for two out of the three datasets. These results are expected due to the fewer frequency components used in the SM-ES trajectory optimization compared to the standard ergodic search optimization. SM-ES also outperforms the Greedy, MCTS, and Prior ES planners in terms of final map error for two out of the three datasets. The Cuprite dataset is the exception for the map error, while the ice favorability index map is the exception for the planning time.

	Planning Time (s)	**MAE**
Synthetic Data		
SM-ES vs. Greedy		✓
SM-ES vs. MCTS		✓
SM-ES vs. Baseline ES	✓	
SM-ES vs. Prior ES	✓	✓
Ice Favorability Index		
SM-ES vs. Greedy		✓
SM-ES vs. MCTS		✓
SM-ES vs. Baseline ES		
SM-ES vs. Prior ES	✓	✓
Cuprite Data		
SM-ES vs. Greedy		
SM-ES vs. MCTS		
SM-ES vs. Baseline ES	✓	
SM-ES vs. Prior ES	✓	

TABLE II: Check marks denote statistical significance at the $\alpha = 0.05$ level of SM-ES compared to the listed informative planner and metric, using the one-sided Wilcoxon signed rank test and corrected for multiple comparisons.

B. *Map Error*

Boxplots of the mean absolute error shown in Figure 5 demonstrate that SM-ES performs the best on the ice favorability index data in terms of this metric, even outperforming baseline ergodic search. Despite the lack of significance for this comparison, the distribution of MAE statistics is tighter for SM-ES than for baseline ergodic search. However, the synthetic maps and Cuprite maps exhibit higher variation in MAE for SM-ES. For the synthetic data, the SM-ES distribution exhibits a variance similar to that of the greedy or MCTS planners, while baseline ergodic search exhibits better performance in this regard.

Further investigation of the map error indicates that the mapping method seems to struggle with the Cuprite data regardless of which planner was used. For this dataset, the full coverage method exhibits a much higher final map error that is similar to the other methods' (see Table I). The boxplots in Figure 5 exhibit this difference across the three datasets: while the distributions for the synthetic maps and the ice favorability index are noticeably different between planners, the distributions for the Cuprite data are nearly the same. Further investigation indicates that certain maps in the Cuprite data are much more difficult for the method to learn and are potentially contributing to the high map error and variation in map error shown in Figure 5(c).

VI. Conclusion

Spectral mixture ergodic search (SM-ES) leverages Gaussian process-based mapping techniques to learn the map while also optimizing an ergodic trajectory over that map. This iterative procedure is accomplished by use of the spectral mixture kernel, which determines the frequencies and weights that best describe the covariance structure of the data and that are then used in ergodic search in place of the usual set of frequencies and weights. This procedure enables the search algorithm to adapt to new information from the map and has been demonstrated on synthetic maps and real-world datasets consisting of an ice favorability index for the lunar south polar region and mineral maps from Cuprite, Nevada. For two out of three datasets, SM-ES outperforms other informative path planners except for baseline ergodic search in terms of map prediction error, while being faster than baseline ergodic search due to the more limited set of frequencies on which it operates.

Acknowledgements This work was supported by the NASA Space Technology Graduate Research Opportunity (80NSSC22K1206), NASA LSITP award 80MSFC20C0008 to Astrobotic Technologies, the AI Institute for Societal Decision Making (NSF 2229881), the Jet Propulsion Laboratory's Strategic University Research Partnerships (SURP) program, and the AI Institute for Resilient Agriculture (NSF 024408B).

REFERENCES

[1] NASA, "Artemis III Science Definition Team Report," NASA, Tech. Rep. NASA/SP-20205009602, 2021.

[2] J. D. Baker, J. O. Elliott, J. T. Keane, N. R. Khan, R. P. Kornfeld, H. D. Nayar, and I. A. Nesnas, "The Endurance Lunar Rover Sample Return Mission," in *2024 IEEE Aerospace Conference*. IEEE, 2024.

[3] K. Di, M.-H. Zhu, Z. Yue, Y. Lin, W. Wan, Z. Liu, S. Gou, B. Liu, M. Peng, Y. Wang, S. Niu, J. Zhang, J. Li, J. Xie, L. Xi, J. Yang, and B. Xue, "Topographic Evolution of Von Kármán Crater Revealed by the Lunar Rover Yutu-2," *Geophysical Research Letters*, vol. 46, 2019.

[4] S. Vijayan, K. B. Kimi, A. Chavan, R. Aditi, U. Thahira, V. R. Subramanian, R. K. Sinha, Amitabh, S. Vadawale, M. Shanmugam, N. P. S. Mithun, A. R. Patel, A. B. Sarbadhikari, K. V. Iyer, K. Suresh, A. Prashar, G. Rima, and A. Bhardwaj, "Chandrayaan-3 landing site evolution by South Pole-Aitken basin and other impact craters," *Icarus*, vol. 425, 2025.

[5] A. Candela, K. Edelson, and D. Wettergreen, "Mars Rover Exploration Combining Remote and In Situ Measurements for Wide-Area Mapping," 2020.

[6] D. R. Thompson and D. Wettergreen, "Intelligent Maps for Autonomous Kilometer-Scale Science Survey," 2008.

[7] B. Shah and H. Choset, "Survey on urban search and rescue robots," *Journal of the Robotics Society of Japan*, vol. 22, no. 5, 2004.

[8] Y. Brouwer, A. Vale, and R. Ventura, "Informative path planner with exploration–exploitation trade-off for radiological surveys in non-convex scenarios," *Robotics and Autonomous Systems*, vol. 136, Feb. 2021.

[9] K. Ennico-Smith, A. Colaprete, D. Lim, and D. Andrews, "The VIPER Mission, a Resource-Mapping Mission on Another Celestial Body," in *Space Resources Roundtable XXII Meeting*, Jun. 2022.

[10] R. Trautner, S. Barber, R. Fisackerly, D. Heather, B. Houdou, C. Howe, S. Iacobellis, M. Leese, A. Mariani, G. Meogrossi *et al.*, "PROSPECT: A comprehensive sample acquisition and analysis package for lunar science and exploration," *Frontiers in Space Technologies*, vol. 5, 2024.

[11] J. Quinn, J. Captain, A. Eichenbaum, R. Aguilar-Ayala, J. Kleinhenz, K. Zacny, P. Chu, and V. Vendiola, "Polar resources ice mining experiment-1 (PRIME-1) NASA's first polar drilling and volatiles detection mission," in *Space Resources Roundtable XXIII Meeting*, 2023.

[12] L. Schweitzer, H. Jamal, H. Jones, D. Wettergreen, and W. L. Red Whittaker, "Micro Rover Mission for Measuring Lunar Polar Ice," in *2021 IEEE Aerospace Conference (50100)*. Big Sky, MT, USA: IEEE, Mar. 2021.

[13] C. Gross, M. Al-Samir, J. L. Bishop, F. Poulet, F. Postberg, and D. Schubert, "Prospecting in-situ resources for future crewed missions to Mars," *Acta Astronautica*, vol. 223, Oct. 2024.

[14] K. I. Alvarado and M. C. Bazzocchi, "Quantification method for assessment of asteroid resource accessibility," in *2022 IEEE Aerospace Conference (AERO)*, 2022.

[15] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4, 2011.

[16] K. M. Cannon and D. T. Britt, "A geologic model for lunar ice deposits at mining scales," *Icarus*, vol. 347, Sep. 2020.

[17] A. Candela, D. Thompson, E. N. Dobrea, and D. Wettergreen, "Planetary robotic exploration driven by science hypotheses for geologic mapping," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep. 2017.

[18] F. Gerber and D. Nychka, "Fast covariance parameter estimation of spatial Gaussian process models using neural networks," *Stat*, vol. 10, 2021.

[19] J. Van Breedam, P. Huybrechts, and M. Crucifix, "A Gaussian process emulator for simulating ice sheet–climate interactions on a multi-million-year timescale: CLISEMv1.0," *Geoscientific Model Development*, vol. 14, Oct. 2021.

[20] L. M. Miller and T. D. Murphey, "Trajectory optimization for continuous ergodic exploration," in *2013 American Control Conference*, Jun. 2013.

[21] Z. Ren, A. K. Srinivasan, H. Coffin, I. Abraham, and H. Choset, "A Local Optimization Framework for Multi-Objective Ergodic Search," in *Robotics: Science and Systems XVIII*. Robotics: Science and Systems Foundation, Jun. 2022.

[22] M. Böhmer and M. Kucera, *Prospecting and exploration of mineral deposits*. Elsevier, 2013.

[23] R. Zuo, "Geodata Science-Based Mineral Prospectivity Mapping: A Review," *Natural Resources Research*, vol. 29, Dec. 2020.

[24] H. Brown, A. Boyd, B. Denevi, M. Henriksen, M. Manheim, M. Robinson, E. Speyerer, and R. Wagner, "Resource potential of lunar permanently shadowed regions," *Icarus*, vol. 377, 2022.

[25] H. Rajesh, "Application of remote sensing and GIS in mineral resource mapping-An overview," *Journal of Mineralogical and Petrological Sciences*, vol. 99, no. 3, 2004.

[26] A. M. Melesse, Q. Weng, P. S. Thenkabail, and G. B. Senay, "Remote sensing sensors and applications in environmental resources mapping and modelling," *Sensors*, vol. 7, no. 12, 2007.

[27] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.

[28] D. E. Dong, H. P. Berger, and I. Abraham, "Time Optimal Ergodic Search," in *Robotics: Science and Systems*, 2023.

[29] Z. Ren, A. K. Srinivasan, B. Vundurthy, I. Abraham, and H. Choset, "A pareto-optimal local optimization framework for multiobjective ergodic search," *IEEE Transactions on Robotics*, vol. 39, no. 5, 2023.

[30] G. Sartoretti, A. Rao, and H. Choset, "Spectral-based distributed ergodic coverage for heterogeneous multi-agent search," in *International Symposium Distributed Autonomous Robotic Systems*. Springer, 2021.

[31] C. Hughes, H. Warren, D. Lee, F. Ramos, and I. Abraham, "Ergodic trajectory optimization on generalized domains using maximum mean discrepancy," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 2025.

[32] M. M. Sun, A. Gaggar, P. Trautman, and T. Murphey, "Fast ergodic search with kernel functions," *IEEE Transactions on Robotics*, vol. 41, 2025.

[33] L. Dressel and M. J. Kochenderfer, "On the Optimality of Ergodic Trajectories for Information Gathering Tasks," in *2018 Annual American Control Conference (ACC)*. Milwaukee, WI, USA: IEEE, Jun. 2018.

[34] A. Rao, A. Breitfeld, A. Candela, B. Jensen, D. Wettergreen, and H. Choset, "Multi-objective ergodic search for dynamic information maps," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023.

[35] A. G. Wilson and R. P. Adams, "Gaussian Process Kernels for Pattern Discovery and Extrapolation," in *Proceedings of the 30th International Conference on Machine Learning*, Dec. 2013.

[36] A. G. Wilson, "Correction to Spectral Mixture (SM) Kernel Derivation for Multidimensional Inputs," May 2015.

[37] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," 2017.

[38] J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson, "GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. Curran Associates Inc., 2018.

[39] H. Landau, "Sampling, data transmission, and the Nyquist rate," *Proceedings of the IEEE*, vol. 55, Oct. 1967.

[40] A. Candela Garza, "Bayesian models for science-driven robotic exploration," Ph.D. dissertation, 2021.

[41] R. F. Kokaly, R. N. Clark, G. A. Swayze, K. E. Livo, T. M. Hoefen, N. C. Pearson, R. A. Wise, W. Benzel, H. A. Lowers, R. L. Driscoll, and A. J. Klein, *USGS Spectral Library Version 7*, 2017.

[42] S. Kodgule, A. Candela, and D. Wettergreen, "Non-myopic Planetary Exploration Combining In Situ and Remote Measurements," in *Proceedings of (IROS) IEEE/RSJ International Conference on Intelligent Robots and Systems*, November 2019.

[43] F. Wilcoxon, "Individual Comparisons by Ranking Methods," *Biometrics Bulletin*, no. 6, Dec 1945.