Resource Limited Exploration and Coverage through Ergodic Optimization

Ananya Rao

CMU-RI-TR-23-11

April 17, 2023

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

Thesis Committee:
Howie Choset, *chair*
David Wettergreen
Ian Abraham
Brady Moon

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

Abstract

Effective exploration and coverage under resource limitations is crucial for many applications such as planetary exploration and search and rescue. Optimizing the use of limited resources while effectively exploring an area is vital in scenarios where sensing is expensive, exhaustive, or has adverse effects. In this thesis we present a novel, sparse sensing motion planning algorithm for autonomous mobile robots in coverage problems with limited sensing resources. We approach this problem using ergodic search processes, which produce trajectories that drive robots to spend time in areas in proportion to the expected amount of information in those regions. We recast the ergodic search problem as a mixed-integer optimization problem in order to determine when and where a sensor measurement should be taken while optimizing the agent's trajectory for coverage. We further employ a continuous relaxation of the presented sparse ergodic optimization problem to reduce computation time. We show that our approach performs comparably to dense sampling methods in terms of coverage performance, by gathering information-rich measurements while adhering to sensing resource constraints. We extend our formulation to problems involving multiple agents, and experiments demonstrate the capability of our approach to automatically distribute sensing resources across a team. Additionally, experiments demonstrate the applicability of our approach to both synthetic and real-world data.

Acknowledgments

I owe many thanks to my advisor, Professor Howie Choset, for his incredible support over the years. His guidance and insight have been indispensable not only in this project, but also in my growth as a researcher.

Further, I would like to thank the other members of my committee: Professor David Wettergreen, Professor Ian Abraham, and Brady Moon. Their feedback and varied perspectives on this work, as well as insights on how it connects to different fields, has been invaluable.

I am incredibly grateful to Professor Guillaume Sartoretti, for his continued mentorship and support throughout my journey in academia.

I also would like to thank my colleagues in the Biorobotics Lab, particularly Ben Freed, for providing engaging and thought-provoking conversations, asking the hard questions about my ideas, and creating an enriching work environment.

Many thanks to the members of the Wettergreen Lab, especially Abigail Breitfeld and Maggie Hansen, for their insights and feedback on this work, and for sharing (and helping understand) data from field testing.

I would like to thank my incredible friends and family for their unconditional support. Thank you to Lori, Amanda, and Benj for being my sanity checks, Swapnil, Abby, Mrinal, and Sam for being sounding boards, and all of my friends for being irreplaceable parts of my support network. And finally, to Amma, Appa, and Akash, thank you for giving me the courage and strength to do any of this.

Contents

List of Figures

Chapter 1

Introduction

1.1 Motivation

Robotic systems frequently need to explore and search for information under resource constraints. For example, data collection and computation may be limited by on-board resources. Optimizing the use of sensing resources in particular is important in applications where sensing is expensive, exhaustive, or can have adverse effects.

Figure 1.1: **Motivational Applications:** A rover analog conducting a science mission field test in Cuprite, NV (left), a theorized use case of drones in agricultural applications [1] (center), and a concept figure of in-orbit satellite inspection (right).

Some applications where sensing resources could be limited include scientific planetary exploration, agricultural robotics, and satellite inspection (depicted in Fig 1.1). In planetary exploration for science missions using a rover, the rover may need to stop in order to take a high-quality sensing measurement, which could in turn shorten the total operating time available during a mission [24]. In some agriculture or forestry applications the total number of available sensing measurements could be limited [16]. For example, if the goal is to take measurements to calculate the carbon capture potential of a vegetated region by placing sensors, the number of available sensors could be limited [26]. For in-orbit satellite inspection, we can gain more information about regions of the satellite through interactive sensing modalities (e.g. probing). However, each interaction with the satellite being inspected risks damaging the satellite, or knocking the satellite out of orbit.

In all of these application areas, the overarching goal is still to cover the target region, though the available sensing resources are limited. Therefore, there is a need for approaches that can effectively cover a given region, while deciding when to use limited resources.

1.2 Challenges

This work considers the scenario where a robot must search for information using sensors that incur a non-trivial cost to activate. We want to develop an approach that will drive robotic agents to intelligently and economically use limited resources, specifically limited sensing resources.

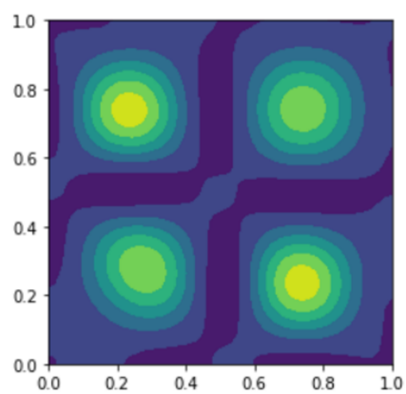

Figure 1.2: **Example of Motion Model Dependency:** An information distribution where moving between regions of high expected information (depicted in yellow) requires driving through regions of low expected information (depicted in purple) for a curved-constrained motion model.

In order to do this, we need to consider the dependency of resource use on both the information that has already been acquired, and also the mobility capabilities of the robot. For example, if there are multiple regions of high expected information (or areas that we want to take measurements in) separated by a region of low expected information (or an area that is not very important to take measurements in), depending on the robot's mobility constraints, it may still need to traverse that regions of low information to move between regions of high information (example shown in Fig 1.2).

1.3 Contributions

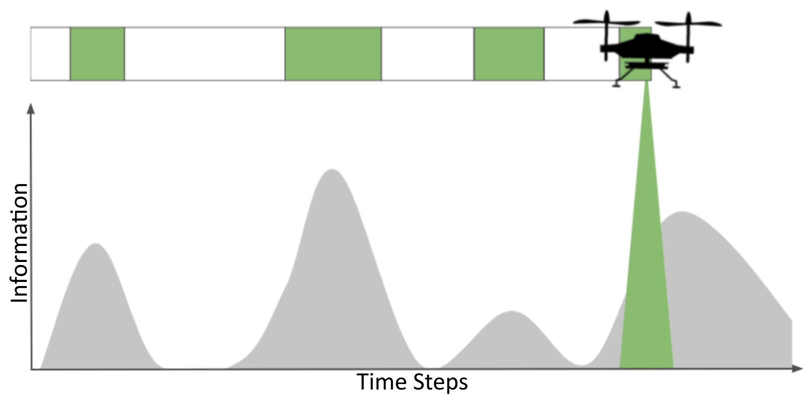

Figure 1.3: **Proposed Approach for Sparse Ergodic Optimization:** Our approach automates how a robot explores an area and what informative measurements to collect in a joint optimization problem. Illustrated above is an example sparse sensing solution for covering a one-dimensional information distribution (gray distribution), where peaks correspond to areas of high expected information, and the green colored bar represents when the robot takes a measurement.

In this work, we hypothesize that explicitly planning for when and where to use limited sensing resources will improve the coverage performance that is possible with fewer sensing measurements. This is based on the hypothesis that there is a required set of informative measurements needed to fully characterize a given coverage problem, and this is different for different coverage problems. In order to account for dependency of resource use on the mobility capabilities of the agents, we propose jointly optimizing the robot's trajectory and when and where the robot should take a sensing measurement. That is, we will consider the robot's trajectory and sensing or sampling locations as a single optimization problem. We formulate this joint optimization approach as an extension of ergodic trajectory optimization.

We pose the problem of acquiring a set of required sensor measurements as a sparse ergodic optimization problem, where the decision to take a measurement is optimized as a vector of decision variables, and show that placing an L^1 norm can promote sparsity in this vector. Using a continuous relaxation of the resulting mixed-integer problem, we show that our formulation performs comparably to dense sampling methods, collecting information-rich measurements while adhering to limited sensing measurements. Further comparisons show comparable performance with the continuous relaxation of the mixed-integer program while reducing computational resources.

This formulation extends to multi-agent path planning problems. For multi-agent teams, our approach automatically distributes sensing resources amongst the team. Our experiments also demonstrate that our approach works on both synthetic information maps, and on real-world data.

Chapter 2

Background

2.1 Coverage Planning Methods

Current coverage planning methods generally fall into one of three main categories: geometric, gradient-based, and trajectory optimization-based approaches. Geometric methods, e.g., lawnmower patterns, can be good search strategies in order to uniformly cover a domain in which there is near-uniform probability of finding a target [3, 12]. Since these approaches exhaustively cover the search domain, they are also the logical choice in cases where there is no *a priori* information about the targets' locations.

An information map, or information distribution, is defined to be a probability distribution representing the likelihood of a target being found at each location in the domain. When such *a priori* information is available (and, usually, non-uniform), more advanced search processes can be created that leverage this information map in order to improve coverage according to some metric, such as the amount of information gained.

For example, in gradient-based, or "information surfing", methods [8, 17, 28], agents guide their movement in the direction of the derivative of the information map around their positions to greedily maximize the short-term information gain. That is, agents are always driven in the direction of the greatest information gain, which naturally leads them to areas where the likelihood of finding a target is maximized. Information surfing can be implemented in a fully decentralized manner, since it does not require tight coordination between agents, and potential fields can be introduced to help distribute agents to different areas of the domain. However, gradient-based approaches generally do not rely on the uncertainty associated with the information distribution, which can lead to areas left unexplored, as this uncertainty can help differentiate areas of low-information that have not been explored from areas with no information to be gained. Gradient-based approaches are also very sensitive to noise in the information map, as the gradient cannot be estimated accurately in these situations, and suffer from greedily over-exploiting local information maxima.

Optimization-based approaches look at search as an information gathering maximization problem, which is then solved by planning (usually joint) paths for the agents. Several recent works in coverage methods [6, 7, 20, 23] rely on sampling-based path planning, where a large number of paths are sampled and the best path is chosen based on a cost metric. Optimization-based approaches can combine both the predicted information distribution as well as its associated uncertainty into the cost function that drives the optimization. However, these approaches generally do not scale well for large multi-agent systems since they remain centralized. Even for sampling-based approaches, the number of paths that need to be sampled to find near-optimal search paths grows exponentially with the number of agents, although growing the number of samples linearly with the team size seems to experimentally provide good-quality search paths [6, 7].

2.2 Limited Sensing Planning Methods

Many real-world robotic applications are plagued by resource limitations. Sparse sensing techniques are useful in these scenarios. Most prior work in applications involving sparse sensing focus on using sparse sensor measurements (and therefore sparse data) to accomplish tasks like localization and SLAM [19], depth reconstruction [18] and wall-following [29]. These works mostly explore methods to better use limited data that has already been acquired, or that is actively being acquired. However, in all of these approaches, the robot still has to acquire the measurements, and post-optimizing for sparse data points does not help reduce costs for limited onboard resources. While intelligently using limited data does help improve the performance of resource-limited robotic systems, further improvements can be made by deciding where to take these limited measurements.

Another perspective on sampling location planning is next best view planning methods. These approaches typically focus on finding a set of good "viewpoints", that is sensing or sampling locations for a problem, often to accomplish tasks like exploration [9] and modeling [10]. Next best view techniques typically optimize for taking informative measurements, and so, they accomplish the goal of intelligently acquiring data or measurements. However, many of these methods do not usually explicitly consider agent dynamics.

2.3 Ergodic Search

Most information-based search and coverage methods view the information gathering problem through one of two lenses: exploration or exploitation. Through the exploratory lens, the information acquisition problem is framed as the wide space search for diffused information, for applications like localization or coverage [4, 22]. On the other hand, through the exploitative lens, the information gathering problem is framed as the direct search for highly structured information peaks, such as in object search [4, 5]. Ergodic search is able to balance both exploration and exploitation goals by accounting for both diffused information densities and highly focused points of information [7, 20, 21]. By doing this, ergodic control trajectories are able to conduct both broad-stroke coverage for diffused information densities and localized search for more focused high information points, thereby balancing exploration and exploitation. Specifically, an ergodic path will drive a robot to spend more time in regions of higher expected information in an *a priori* information map, and less time in regions of low information.

Ergodic coverage [20] produces trajectories for agents such that they spend time in areas of the domain proportional to the expected amount of information present in that area. In order to control agents to accomplish this behavior, we pose an optimization problem that minimizes the distance between the time-average statistics of the agent Eq 2.1 and the underlying information map.

$$C^t(\mathbf{x}, \gamma_t) = \frac{1}{t}\sum_{\tau=0}^{t-1} \delta(\mathbf{x} - \gamma_i(\tau)), \tag{2.1}$$

where γ is the agent's trajectory, defined as $\gamma : (0, t] \rightarrow \mathcal{X}$, t is the discrete time horizon, and δ is the Dirac delta function, with $\mathcal{X} \subset \mathbb{R}^{\mathrm{d}}$ in the d-dimensional search domain. The spatial time-average statistics of an agent's trajectory quantifies the fraction of time spent at a position $\mathbf{x} \in \mathcal{X}$.

The expected information distribution, or information map, over the domain to be explored and searched is determined by a target distribution which defines the likelihood of generating informative measurements at any given location in the search domain. Formally, the agent's time-averaged trajectory statistics are optimized against this expected information distribution over the whole domain, by minimizing the distance between the Fourier spectral decomposition of each distribution. This is obtained by minimizing the ergodic metric $\Phi(\cdot)$, expressed as the weighted sum of the difference between the spectral coefficients of these two distributions [20]:

$$\Phi(\gamma(t)) = \sum_{k=0}^{m} \alpha_k \left| c_k(\gamma_t) - \xi_k \right|^2 , \tag{2.2}$$

where c_k and ξ_k are the Fourier coefficients of the time-average statistics of an agent's trajectory $\gamma(t)$ and the desired spatial distribution respectively, and α_k are the weights of each coefficient difference. In practice, $\alpha_k = \sqrt{(1 + \|k\|^2)^{-(d+1)}}$ is usually defined to place higher weights on lower frequency components, which correspond to larger spatial-scale variations in the information distribution.

The goal of ergodic coverage is to generate optimal controls $\mathbf{u}^*(t)$ for an agent, whose dynamics are described by a function $f : \mathcal{Q} \times \mathcal{U} \to \mathcal{T}\mathcal{Q}$, such that

$$\begin{aligned} \mathbf{u}^*(t) \quad &= \quad \operatorname{argmin}_{\mathbf{u}} \Phi(\gamma(t)), \\ \text{subject to } \dot{\mathbf{q}} \quad &= \quad f(\mathbf{q}(t), \mathbf{u}(t)), \;\; \|\mathbf{u}(t)\| \leq u_{max} \end{aligned} \tag{2.3}$$

where $\mathbf{q} \in \mathcal{Q}$ is the state and $\mathbf{u} \in \mathcal{U}$ denotes the set of controls. Eq.2.3 is solved directly for the optimal control input at each time-step, by trajectory optimization to plan feed-forward trajectories over a specified time horizon [21], or by using sampling-based motion planners [15], where it is straightforward to pose additional constraints such as obstacle avoidance. Note that in this optimization, each point $x(t)$ in the robot's trajectory is considered a sample point during the exploration of the search domain, which can be detrimental in scenarios where the robot's trajectory traverses over large regions of low expected information. In the following section, we describe our approach to selecting the sampling decision times as a part of the optimization.

Chapter 3

Sparse Ergodic Optimization

3.1 Sparse Ergodic Optimization Problem Formulation

In many search and coverage problems, a model of the region can be comprehended and optimized using only a few data points, or informative sensor measurements. However, in many trajectory optimization approaches the locations at which sensing measurements are taken are temporally uniform. In ergodic optimization, this means that each measurement along the trajectory is treated as being equally important, and contributes equally to the ergodic metric (Eq 2.2). Due to both the dynamic constraints of the robots being used for coverage, as well as the non-uniform distribution of information in the region being covered, this can lead to extraneous measurements that do not contribute to modeling the information prior.

For example, in a scenario where there are few areas of high information, separated by large areas of low expected information, uniform sensing measurements would result in many measurements that provide low information gain (see Fig. 3.1). Here, coverage performance, in terms of the ergodic metric, would be improved by mostly taking sensor measurements near the areas of higher information.

Ergodic optimization can be reformulated in order to find this set of optimal measurements for a given search scenario. We optimize for the set of optimal sensor

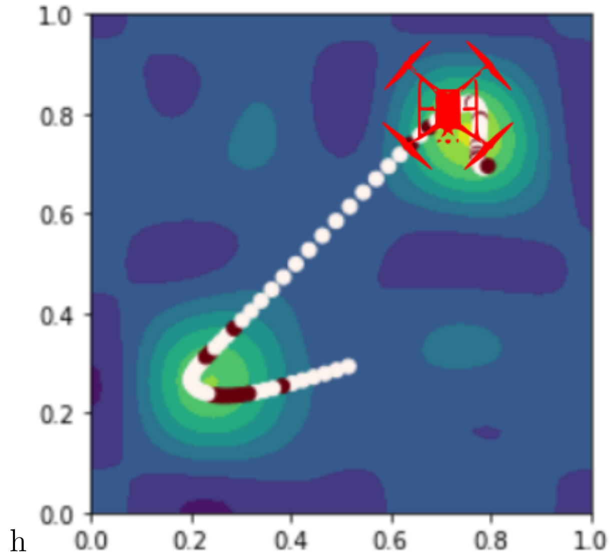

Figure 3.1: **Sparse Ergodic Trajectory Example:** Trajectories are generated with a decision variable for when to take a sensor measurement. This approach results in trajectories which provide maximal information gathering with minimal sensing resources. Points in red indicate sampling locations, while points in white indicate parts of the trajectory where no measurements are taken.

measurements by posing the following ergodic optimization,

$$\begin{aligned} \mathbf{u}^*(t), \lambda^*(t) &= \operatorname{argmin}_{\mathbf{u},\lambda} \Phi_{\text{sparse}}(\gamma(t)), \\ \text{subject to } \dot{\mathbf{q}} &= f(\mathbf{q}(t), \mathbf{u}(t)), \ \ \|\mathbf{u}(t)\| \leq u_{max} \end{aligned} \tag{3.1}$$

where $\mathbf{q} \in \mathcal{Q}$ is the state, $\mathbf{u} \in \mathcal{U}$ denotes the set of controls, and $\lambda(t) \in \{0, 1\}$. $\lambda(t)$ represents the decision variable for choosing whether to take a sensor measurement or not at a given location in the search domain. We promote sparsity in the sample measurements by regularizing λ with an $L1$ optimization [25].

In order to jointly optimize the agent's trajectory and sensing measurement location, we augment the ergodic metric in Eq. 2.2 in the following manner

$$\Phi_{\text{sparse}}(\gamma(t)) = \sum_{k=0}^{m} \alpha_k \left| c_k(\gamma(t), \lambda(t)) - \xi_k \right|^2 + \sum |\lambda_k|, \tag{3.2}$$

where c_k and ξ_k are the Fourier coefficients of the time-average statistics of the set of agent's trajectories $\gamma(t)$ and the desired spatial distribution of agents respectively, and α_k are the weights of each coefficient difference.

The spatial time-average statistics of the agent's trajectory (from Eq. 2.1) are also modified to be,

$$C'^t(\mathbf{x}, \gamma(t)) = \frac{1}{\sum_t \lambda(t)} \sum_{\tau=0}^{t} \lambda(t)\delta(\mathbf{x} - \gamma(\tau)), \tag{3.3}$$

where $\lambda(t) \in \{0, 1\}$.

3.2 Sparse Ergodic Optimization Formulation for Multiple Agents

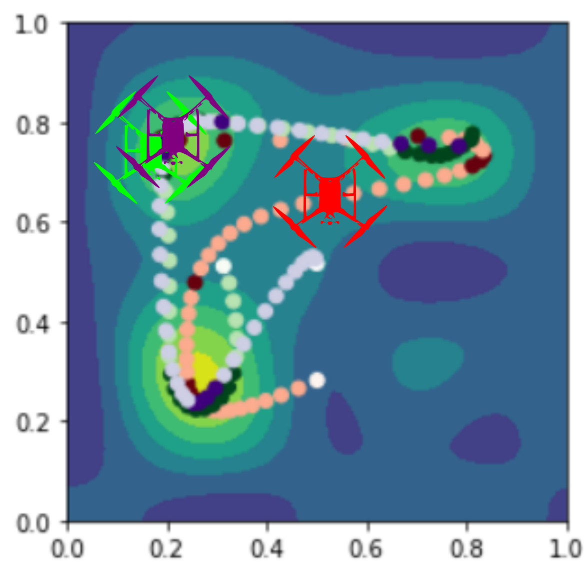

Figure 3.2: **Multi-Agent Sparse Ergodic Trajectory Example:** Trajectories for a multi-agent team are generated with a decision variable for when to take a sensor measurement. This approach automatically updates trajectories which provide maximal information gathering with minimal sensing resources. Darker points indicate sampling locations, while lighter points indicate parts of the trajectories where no measurements are taken.

For a multi-agent team covering a given information prior, the limited number of measurements required to fully cover the region can be distributed among the different agents. The sparse ergodic optimization formulation described in Section 3.1 can be extended to multi-agent settings. For a multi-agent coverage problem, our approach naturally and automatically distributes limited sensor measurements across the team of robotic agents. A sample result is shown in Fig 3.2.

For N agents, we define a set Λ of N decision variable vectors $\lambda_i, i \in [0, N)$ where $\lambda_i(t)$ represents the decision variable for choosing whether to take a sensor measurement or not for agent i, at the agent's location in the search domain at time t. We promote sparsity in sensor measurements by regularizing Λ with an L^1 optimization.

For N agents, the modified joint spatial time-average statistics of the set of agent trajectories $\{\gamma_i\}_{i=1}^N$ are defined as

$$C'^t(\mathbf{x}, \gamma(t)) = \frac{1}{Nt \sum_{i=1}^N \sum_t \lambda_i(t)} \sum_0^t \lambda_i(t) \delta(\mathbf{x} - \gamma_i(\tau)), \tag{3.4}$$

where $\lambda_i(t) \in \{0, 1\}$ for all integers $i \in [0, N)$.

The joint optimization of the set of N agent trajectories and sensing decision variables is driven by the multi-agent formulation of the augmented ergodic metric, defined as,

$$\Phi_{\text{sparse}}(\Gamma(t)) = \sum_{k=0}^m \alpha_k \left| c_k(\Gamma(t), \Lambda(t)) - \xi_k \right|^2 + \sum |\Lambda_k|, \tag{3.5}$$

where c_k and ξ_k are the Fourier coefficients of the time-average statistics of the set of agents' trajectories $\Gamma(t)$ and the desired spatial distribution of agents respectively, and α_k are the weights of each coefficient difference.

The optimal controls and set of optimal sensor measurements for each agent are found using the ergodic optimization problem,

$$\begin{aligned} \mathbf{u}^*(t), \lambda^*(t) \quad &= \quad \text{argmin}_{\mathbf{u},\lambda} \, \Phi_{\text{sparse}}(\gamma(t)), \\ \text{subject to } \dot{\mathbf{q}} \quad &= \quad f(\mathbf{q}(t), \mathbf{u}(t)), \;\; \|\mathbf{u}(t)\| \le u_{max} \end{aligned} \tag{3.6}$$

where $\mathbf{q} \in \mathcal{Q}$ is the state, $\mathbf{u} \in \mathcal{U}$ denotes the set of controls, and $\Lambda(t) \in \{0, 1\}$.

3.3 Mixed-Integer Program Relaxation

Thus far, $\lambda(t)$ is defined to be an integer (i.e. $\lambda(t) \in \{0, 1\}$), resulting in Eq 3.6 being a mixed integer programming problem. However, such mixed integer programming problems are difficult to solve, due to a lack of gradient information from the integer variables [27], and due to requiring direct search methods that do not scale with longer time horizons. Additionally, mixed integer programming problems are expensive to solve because the integer variables introduce a large number of additional constraints on the optimization problem. As a result, the computation and memory requirements for solving the optimization problem can increase exponentially with increase in the number of integer variables.

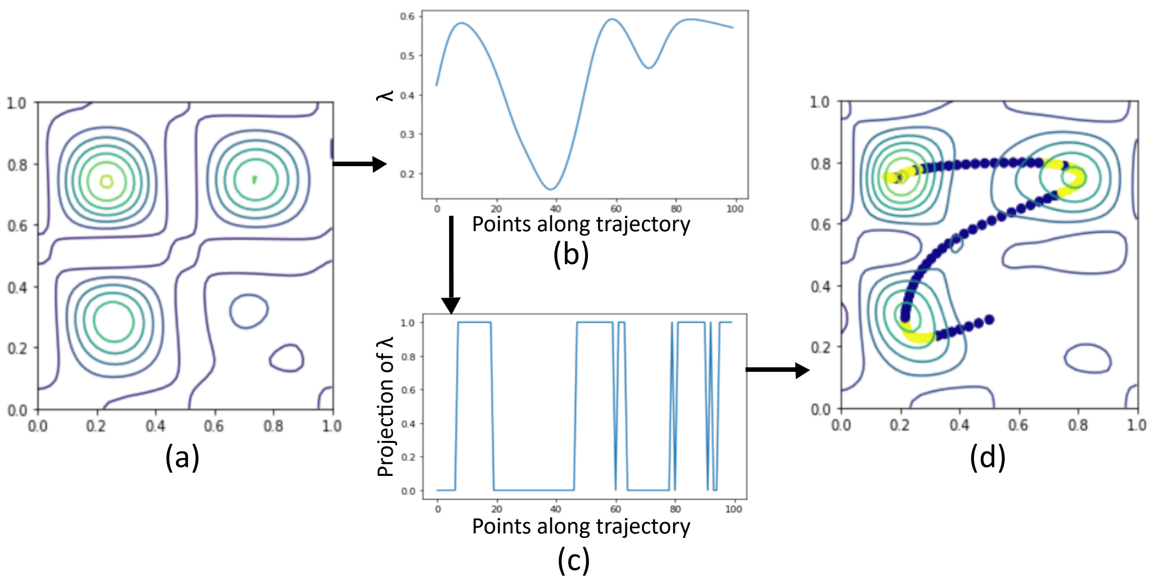

Figure 3.3: **Mixed-Integer Continuous Relaxation for Sparse Ergodic Optimization:** Illustrated is the process with which we jointly optimize for trajectories and when to sample. The decision variable λ is relaxed to be continuous $[0, 1]$ where solutions are projected into the integer $0, 1$ space. As input, our approach takes information prior distributions which guide the planning and sensing. Output trajectory solutions (shown on the right) show concentrated samples over areas of high information with minimal use of sensor resources.

For these reasons, we employ a relaxation of the problem Eq 3.6 by defining $\lambda(t)$ to be a bounded continuous variable $\lambda(t) \in [0, 1]$ and optimize over the new domain.

After optimization, we project λ from the continuous domain to the nearest integer value, while adhering to the sensing budget. This allows us to continuously optimize Eq 3.6, and then map the resultant continuous values of $\lambda(t)$ to discrete values $\{0, 1\}$. The procedure for optimizing the sparse ergodic problem Eq 3.6 is depicted in Fig 3.3.

For the multi-agent case, we employ a relaxation of the problem described by Eq 3.6 by defining each $\lambda_i(t)$ to be a bounded continuous variable $\lambda_i(t) \in [0, 1]$ for all integers $i \in [0, N)$ and optimize over the new domain. After optimization, we project the λ_i for each agent i for all integers $i \in [0, N)$ from the continuous domain to the nearest integer value, while adhering to the sensing budget. We take into account all of the agent trajectories when mapping from the continuous domain to discrete integers in order to distribute sensing measurements across all of the agents.

When we investigate the numerical values of the decision variables $\lambda(t)$ calculated in sparse ergodic optimization Eq. 3.6, we see that there are peaks formed that correspond to peaks in information in the *a priori* information distribution (Fig 3.3). When these continuous numerical values are mapped back to $\lambda(t) \in \{0, 1\}$, this results in $\lambda(t) = 1$ being more likely in areas of higher information, and $\lambda(t) = 0$ being more likely in areas of lower information. This shows that the sparse ergodic optimization drives the likelihood of taking a sensor measurement in an area of higher expected information to be higher, which follows intuition. Additionally, we also see some peaks in regions of lower information, which typically correspond with areas between relatively closely places areas of high information, or with regions of low information where no sensor measurements have been taken for a considerable amount of time.

Chapter 4

Experiment Details

4.1 Data

We evaluate our sparse ergodic optimization formulation on two different kinds of data: synthetic Gaussian information distributions (representing information priors in general search and coverage tasks), and information maps from real data collected during field tests with a planetary rover analog in Cuprite, NV. This section details how the information maps are generated, and how coverage performance on each information map is evaluated.

4.1.1 Synthetic Data

The information maps used for our synthetic data experiments are generated by placing Gaussian peaks at different locations to represent the uncertainty of the corresponding region, with higher values corresponding to higher uncertainty of information. We generate both handcrafted and random Gaussian information maps. We evaluate the coverage of these maps using the ergodic metric described in Section 2.3. A lower ergodic metric signals better coverage of the map. Some example maps are shown in Fig 4.1.

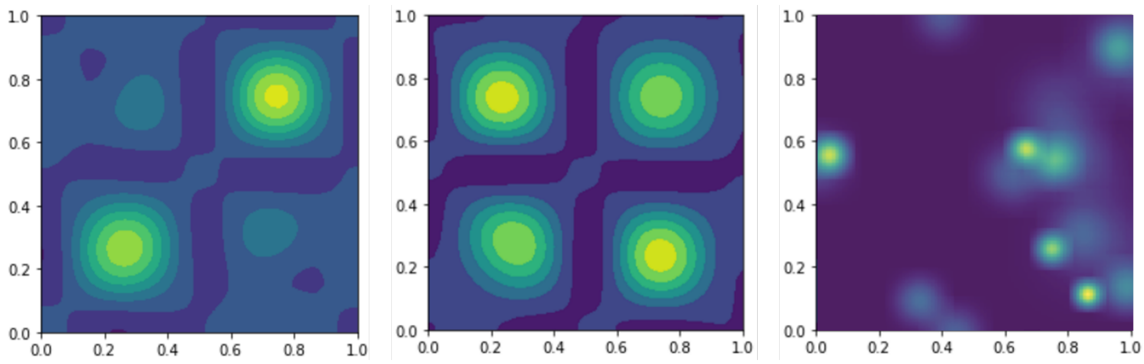

Figure 4.1: **Synthetic Data Information Maps:** Example synthetic Gaussian information maps.

4.1.2 Real-World Data

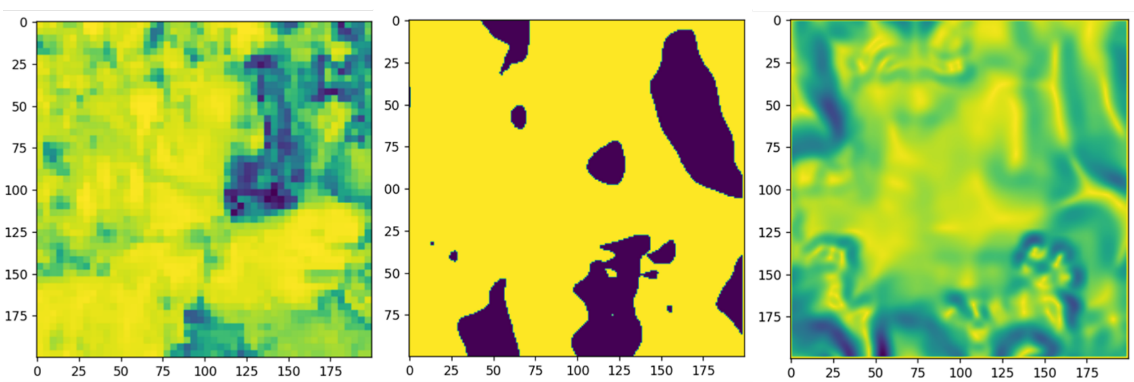

Figure 4.2: **Real-World Data Information Maps:** Example information maps of the entropy map (left), shade map (middle), and slope map (right).

Entropy Map

This information prior in our planetary rover analog experiments is related to the objective of reducing uncertainty in scientific information. We use the entropy map formulation proposed by Candela *et al.*, where "entropy" of each point in the map is related to the variance of a Gaussian process which predicts measurements at each point [11]. We use low-resolution Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) satellite data as the prior and use high-resolution Airborne Visible Near Infrared Spectrometer - New Generation (AVIRIS-NG) data as a proxy for *in-situ* samples [2, 13, 14]. In order to focus the ergodic search on areas of high entropy, we threshold the entropy maps (setting areas of high entropy above 75% of the maximum value to 1, and areas of low entropy below 75% of the maximum value to 0). We find experimentally that thresholding the entropy maps results in better coverage of the region when using an ergodic method. An example of the entropy map is shown in Fig 4.2. We evaluate the coverage of the entropy model using the ergodic metric, where a lower ergodic value indicates better coverage.

Shade Map

Another objective in our planetary rover analog experiments is to explore shadowed regions. Using a digital elevation model (DEM) of the field site, we use raycasting based on the angle of the sun to generate a map of shaded regions. Shadows have a high value while sunlit regions have a low value, which encourages the rover to stay in shaded areas. An example of the shade map is shown in Fig 4.2. To evaluate coverage performance on the shade maps, we again use the ergodic metric.

Slope Map

Another information map that can be modeled in our planetary rover analog experiments is the slope map of the region. The slope of the terrain in our test site acts as a proxy for risk, so building an accurate model of the slope in a region is important. To generate the prior for the risk map, we use a Sobel image filter on a DEM of the region. We opt to use slope as an estimate of risk for simplicity and because of the limited information available for slip characterization [11]. An example of the slope map is shown in Fig 4.2. To evaluate coverage performance, we use the ergodic metric.

4.2 Agent Sensing and Motion Models

In our set of systematic experiments, the dynamics of the agent considered are defined as the constrained discrete time dynamical system

$$x(t+1) = f(x,u) = x(t) + \tanh(u) \tag{4.1}$$

where $x(t) \in \mathbb{R}^2$ and $u \in \mathbb{R}^2$ is constrained by the tanh function bounded $\in [0, 1]$.

The sensor footprint of the agent is modeled as a Gaussian distribution centered at the agent's position, whose variance prescribes a circular observation range $\rho > 0$. The information maps are built using information distributions that are distributed within the continuous search space $X \in [0, L]^2 \subset \mathbb{R}^2$, which is defined as

$$p(x) = \sum_{i=1}^{3} \eta_i \exp(||x - c_i||^2_{\sum_i^{-1}}) \tag{4.2}$$

where $p(x) : X \to \mathbb{R}^+$, and $\eta_i, c_i, \sum_i$ are the normalizing factor, the Gaussian center, and the Gaussian variance respectively.

4.3 Baseline Methods

We compare our sparse ergodic optimization formulation to two different baseline approaches. The first baseline method is standard ergodic optimization as expressed in Eq. 2.2, where sensor measurements are uniformly distributed along the optimized trajectory. An example trajectory generated from this baseline approach is shown in Fig 4.3 The second baseline is a probabilistic heuristic with a two step process: first we optimize an ergodic trajectory, then we sample measurement locations from the distribution of information under the optimized trajectory. Specifically, sensor measurements are taken at parts of the trajectory that pass over regions of high expected information. This baseline method decouples trajectory or path optimization and determining when and where to take sensor measurements. An example trajectory generated from this baseline approach is shown in Fig 4.4

For the multi-agent experiments, we use the same two baseline methods. In this case, the ergodic trajectories of all N agents are jointly optimized. For the first baseline method, sensing measurements are uniformly distributed both between agents, and then also along each agent's optimized trajectory. For the second baseline method, sensing measurement locations are sampled from the distribution of information under all of the agents' trajectories.

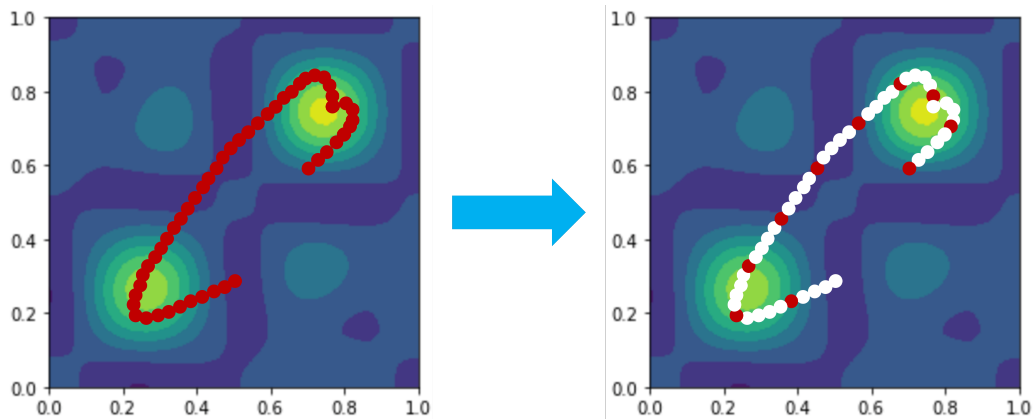

Figure 4.3: **Uniform Sampling Baseline Approach:** Example trajectory generated by the uniform sampling approach, where an ergodic trajectory is found, and then 10 sampling locations are uniformly distributed along the trajectory. Points in red indicate sampling locations, while points in white indicate parts of the trajectory where no measurements are taken.

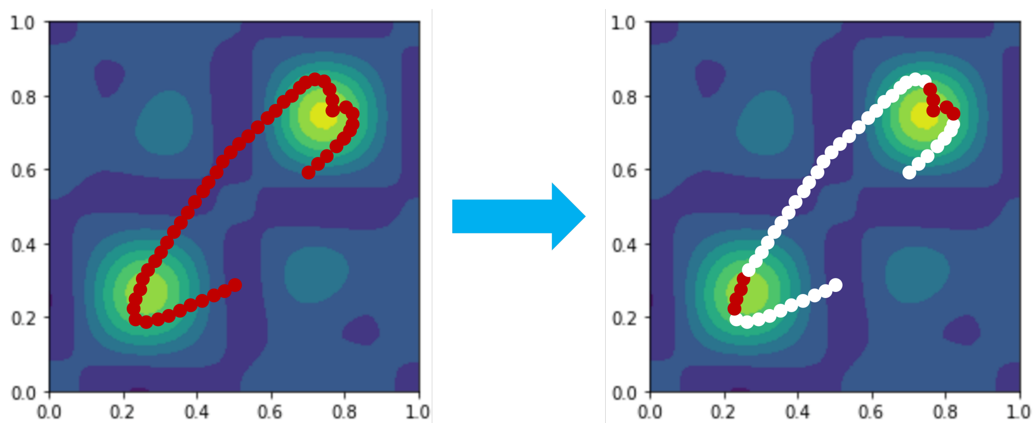

Figure 4.4: **Probabilistic Heuristic Baseline Approach:** Example trajectory generated by the probabilistic heuristic approach, where an ergodic trajectory is found, and then 10 sampling locations are determined from the underlying information distribution. Points in red indicate sampling locations, while points in white indicate parts of the trajectory where no measurements are taken.

4.4 Experiment Setup Details

We investigate the performance of the sparse ergodic optimization described in Section 3.1 in terms of coverage performance, using ergodicity as the metric.

The performance statistics for each method and sensing budget are averaged across 25 randomized experiment setups each, where initial information map is varied between experiments. For each method, 10 different sensing budgets were explored. Agents starting positions, initial information maps, and sensing budgets are kept identical among experiments with different controllers to ensure that our results are comparable.

Further, we investigate the performance of multi-agent sparse ergodic optimization described in Section 3.2 in terms of coverage performance, with ergodicity as the metric. We compare our results to the two baselines described above, averaged across different team sizes (ranging from 3 to 10 agents).

Chapter 5

Results and Discussion

5.1 Single Agent Results

When looking at results of sparse ergodic optimization in terms of overall coverage performance, measured by the ergodic metric, we observe that there is a minimal number of sensor measurements to be taken to minimize ergodicity (see Fig 5.1). In our experiments, this number of sensor measurements varies with changes in information map being covered, sample rate, time horizon and initial sample weights. However, for a set of fixed experiment hyperparameters, the minimal number of samples required is consistent. For any fixed experiment setup, when we take fewer than this minimal number of sensor measurements, the optimization lacks relevant information, and so, we see a decrease in coverage performance. On the other hand, when we take more sensor measurements than the minimal required number, the ergodic value increases, due to extraneous measurements negatively impacting coverage. This is because as

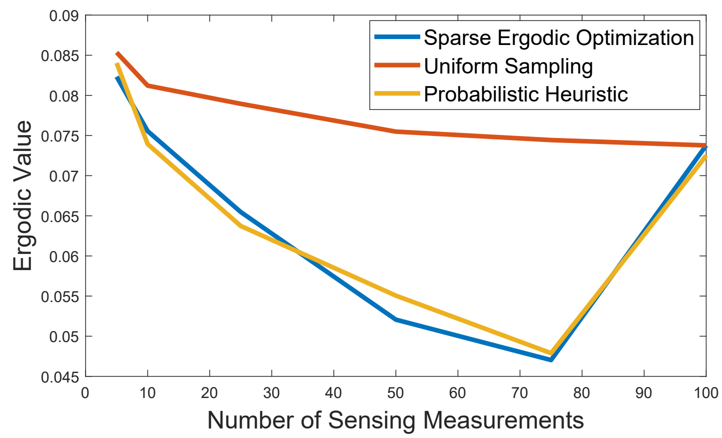

Figure 5.1: **Single Agent Experimental Results on Synthetic Data:** Single agent experiment results show that sparse ergodic optimization has better coverage performance in terms of the ergodic metric when compared to standard ergodic optimization with uniformly distributed sparse measurements. The probabilistic heuristic results in comparable performance.

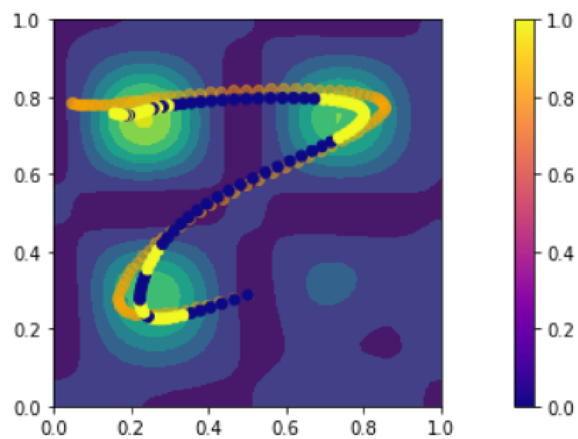

Figure 5.2: **Standard Ergodic and Sparse Ergodic Trajectory Comparison:** Standard ergodic optimization results in a different final trajectory (orange trajectory) than jointly optimizing for both trajectory and sensor weight (blue and yellow trajectory). This implies a cross-effect between agent position and if a sensing measurement is taken, due to which joint optimization becomes more important as the coverage problem scales.

We also experimentally demonstrate that it is better to selectively choose measurements along the standard ergodic trajectory, since in all cases, the coverage performance of the probabilistic heuristic is much better than that of standard ergodic optimization with uniformly distributed measurements. On the other hand, using the probabilistic heuristic on a standard ergodic trajectory has very similar coverage performance to sparse ergodic optimization for a single agent. We see that optimizing for trajectory alone, and jointly optimizing for trajectory and measurement placement creates different resultant trajectories (Fig 5.2), implying that trajectory optimization and measurement choice impact each other. Jointly optimizing for trajectory and sensing measurements leads to lower control cost, as you are directly taking into account the cost of moving between chosen sensing measurements. For a single agent in the simple coverage problems being considered, there aren't large differences in control cost, leading to very similar performance of the probabilistic heuristic and sparse ergodic optimization.

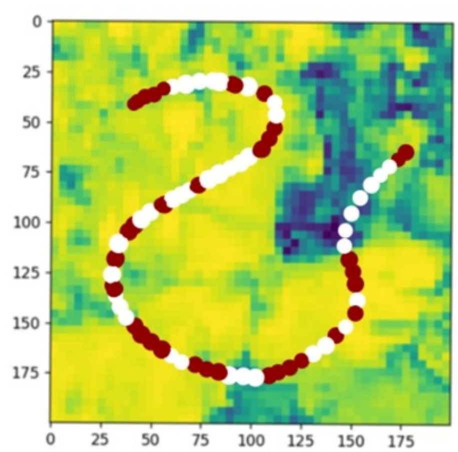

Figure 5.3: **Single Agent Trajectory on Entropy Map:** Example sparse ergodic trajectory over an entropy map. Red points indicate sensor measurement locations, while white points indicate parts of the trajectory where no sensor measurements are taken.

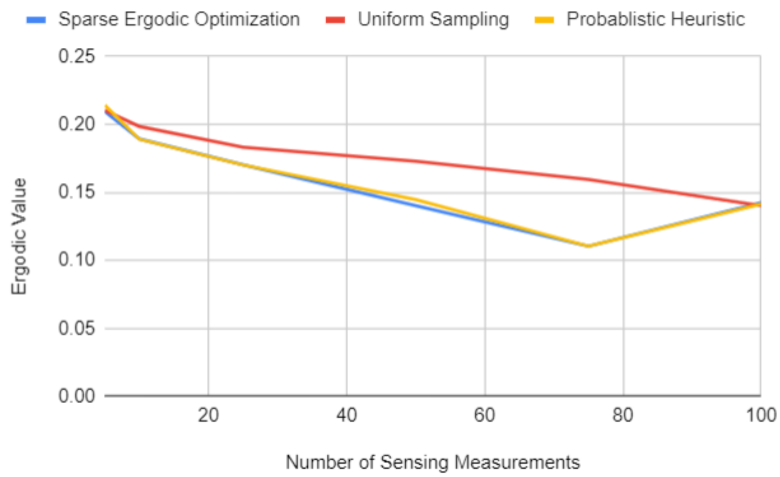

Figure 5.4: **Single Agent Experimental Results on Entropy Maps:** Results show that sparse ergodic optimization has better coverage performance in terms of the ergodic metric when compared to standard ergodic optimization with uniformly distributed sparse measurements. The probabilistic heuristic results in comparable performance.

We observe similar trends in the ergodic values of trajectories from our single agent experiments on the entropy maps (see Fig 5.4). Both sparse ergodic optimization and the probabilistic heuristic baseline method outperform the uniform sampling method. This supports our hypothesis that intelligently choosing sensing locations improves the informativeness of the measurements taken. Further, we again observe that there is a minimal number of sensor measurements to be taken to minimize ergodicity (see Fig 5.4). Note that the number of required sensor measurements for the entropy map is different than that for the synthetic maps, supporting the idea that there is a different set of required measurements in order to fully characterize a given coverage problem. An example of a sparse ergodic trajectory over an entropy map is shown in Fig 5.3.

For single agent experiments on the shade and slope maps, the ergodic values of the resulting trajectories show the same trends as above, with the sparse ergodic optimization and the probabilistic heuristic baseline method performing better than the uniform sampling method (see Fig 5.7 and Fig 5.8). Again we observe that the minimum ergodicity is achieved at different numbers of sensor measurements, supporting the hypothesis that there is a minimum required set of sensor measurements used to fully characterize a coverage problem. For information maps that have a more uniform distribution of areas of interest, like the shade maps, this difference in performance is smaller than the difference in coverage performance seen for information maps with distributed peaks of high expected information. An example of a sparse ergodic trajectory over an shade map is shown in Fig 5.5, and over a slope map in Fig 5.6.

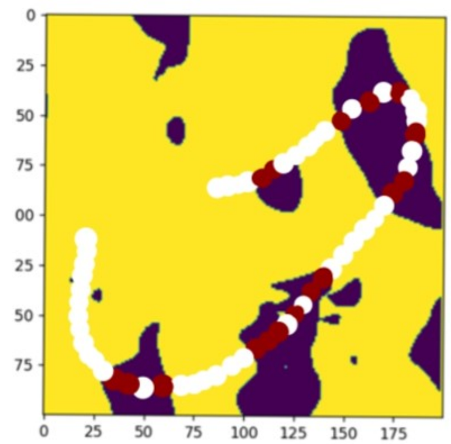

Figure 5.5: **Single Agent Trajectory on Shade Map:** Example sparse ergodic trajectory over a shade map. Red points indicate sensor measurement locations, while white points indicate parts of the trajectory where no sensor measurements are taken.

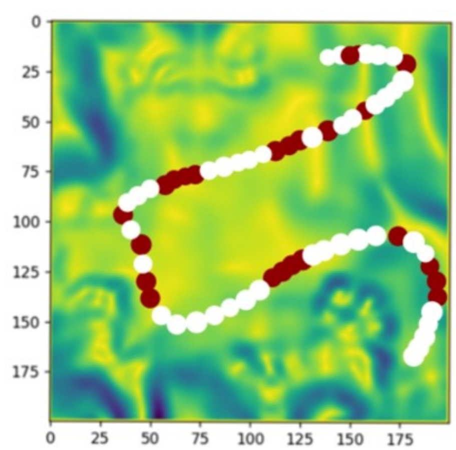

Figure 5.6: **Single Agent Trajectory on Slope Map:** Example sparse ergodic trajectory over a slope map. Red points indicate sensor measurement locations, while white points indicate parts of the trajectory where no sensor measurements are taken.

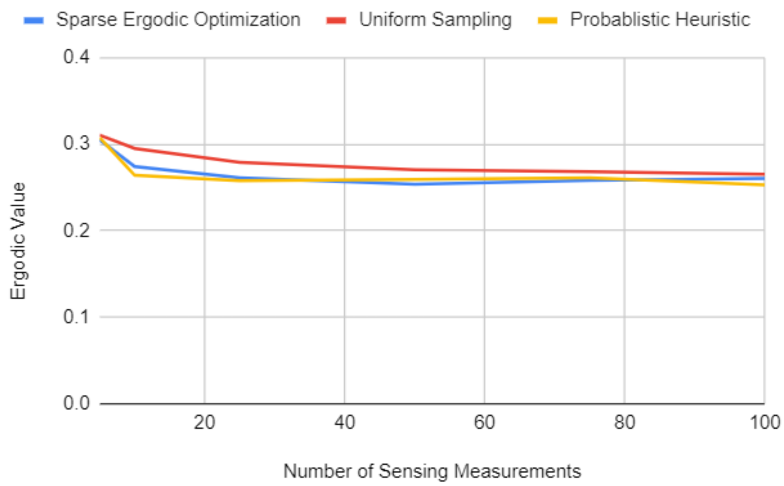

Figure 5.7: **Single Agent Experimental Results on Shade Maps:** Results show that sparse ergodic optimization has better coverage performance in terms of the ergodic metric when compared to standard ergodic optimization with uniformly distributed sparse measurements. The probabilistic heuristic results in comparable performance.

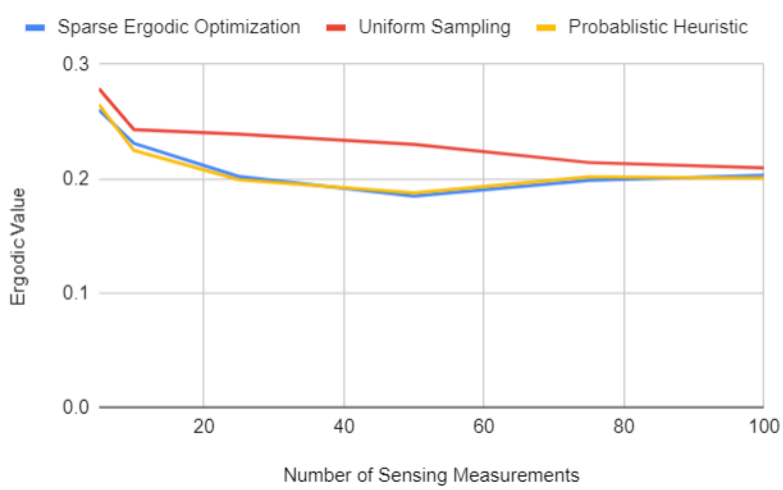

Figure 5.8: **Single Agent Experimental Results on Slope Maps:** Results show that sparse ergodic optimization has better coverage performance in terms of the ergodic metric when compared to standard ergodic optimization with uniformly distributed sparse measurements. The probabilistic heuristic results in comparable performance.

5.2 Multi-Agent Results

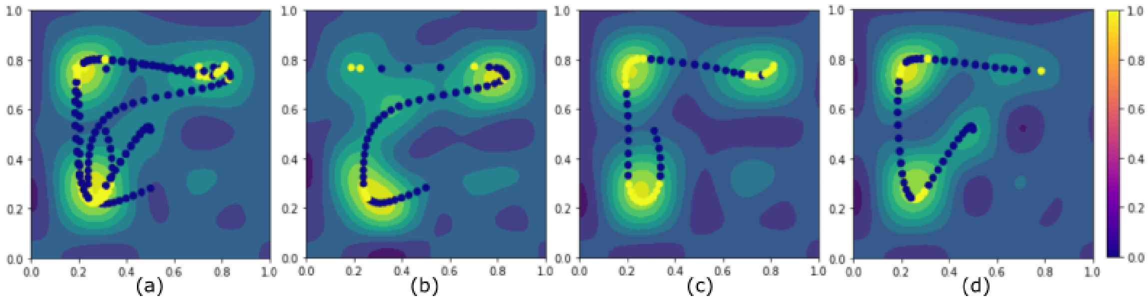

Figure 5.9: **Coverage using Multi-Agent Sparse Ergodic Trajectory:** The sparse sensing trajectories being followed by each of the three agents in the multi-agent team are overlaid on the reconstruction of the information prior using measurements from all agents (a). The yellow points correspond to where a sensing measurement is being taken. The reconstructions of the information prior using the sensing measurements taken by each agent are plotted with the corresponding agent trajectory for agents one (b), two (c), and three (d).

When optimizing limited sensing resources (specifically a sensing budget in terms of a restricted number of sensing measurements) for a multi-agent team using the sparse ergodic optimization approach, we see that the 'workload' of covering different peaks in a given *a priori* information map is distributed among the agents, and the sensing measurements are distributed in order to support the requirements of each coverage workload. An example of this is shown in Fig 5.9.

Similar to single agent sparse ergodic optimization results, we see that using sparse ergodic optimization to distribute a limited sensing budget across a multi-agent team results in improved coverage performance in terms of the ergodic metric with fewer sensing measurements (i.e. for lower sensing budgets) (see Fig 5.10). We also observe that there is a minimal number of sensor measurements that are required in order to minimize ergodicity, and this minimal number varies with changes in experiment hyperparameters like information map being covered, sample rate, time horizon and initial sample weights. For a fixed set of experiment hyperparameters, the ergodicity increases when the sensing budget is increased past this minimum required number, since there are extraneous measurements being taken, while the ergodicity increases with decrease in sensing budget below the minimum required

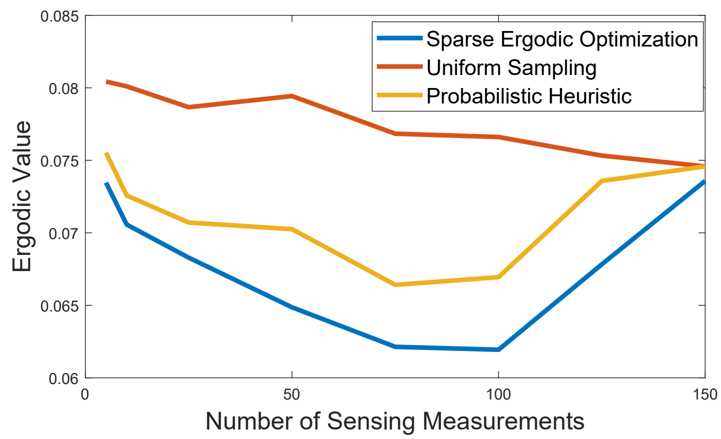

Figure 5.10: **Multi-Agent Experimental Results on Synthetic Data:** Results show that multi-agent sparse ergodic optimization has better coverage performance in terms of the ergodic metric when compared to standard ergodic optimization with uniformly distributed sparse measurements and with sparse measurements distributed using a probabilistic heuristic.

number, since the optimization is missing information.

For multi-agent teams we see that the probabilistic heuristic has worse coverage performance compared to sparse ergodic optimization (see Fig 5.10). As explained in in Section 5.1, jointly optimizing for trajectory and sensing measurements leads to lower control cost. As we scale up the optimization problem, control cost becomes more substantial, as we need to optimize multiple trajectories. Further, the cross-effect of the agents' trajectories and if they take sensing measurements increased for multiple agents, since coverage performance is now impacted by several trajectories with potential areas of overlap. Thus, it becomes more necessary to jointly optimize for trajectory and choosing sensing measurements in multi-agent sparse sensing coverage problems.

We observe similar trends in results from the multi-agent experiments on real-world data (see Fig 5.11, Fig 5.12, and Fig 5.13). In all of these comparisons, both sparse ergodic optimization and the probabilistic heuristic baseline method outperform the uniform sampling baseline method, supporting the idea that intelligently choosing sensing locations improves the informativeness of the measurements taken. For information maps that have a more uniform distribution of areas of interest, like the shade maps, this difference in performance is smaller than the difference in coverage performance seen for information maps with distributed peaks of high expected information.

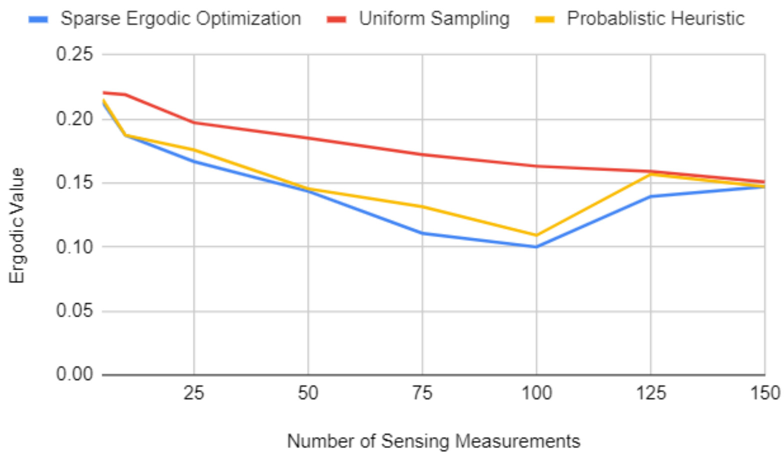

Figure 5.11: **Multi-Agent Experimental Results on Entropy Maps:** Results show that sparse ergodic optimization has better coverage performance in terms of the ergodic metric when compared to standard ergodic optimization with uniformly distributed sparse measurements. The probabilistic heuristic results in comparable performance.

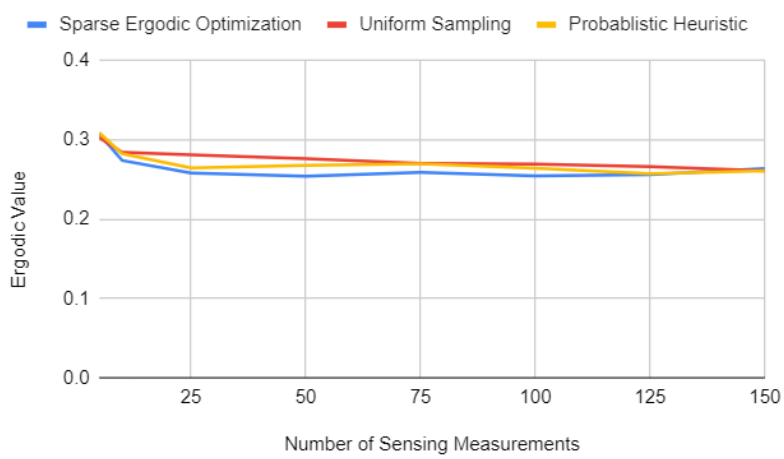

Figure 5.12: **Multi-Agent Experimental Results on Shade Maps:** Results show that sparse ergodic optimization has better coverage performance in terms of the ergodic metric when compared to standard ergodic optimization with uniformly distributed sparse measurements. The probabilistic heuristic results in comparable performance.

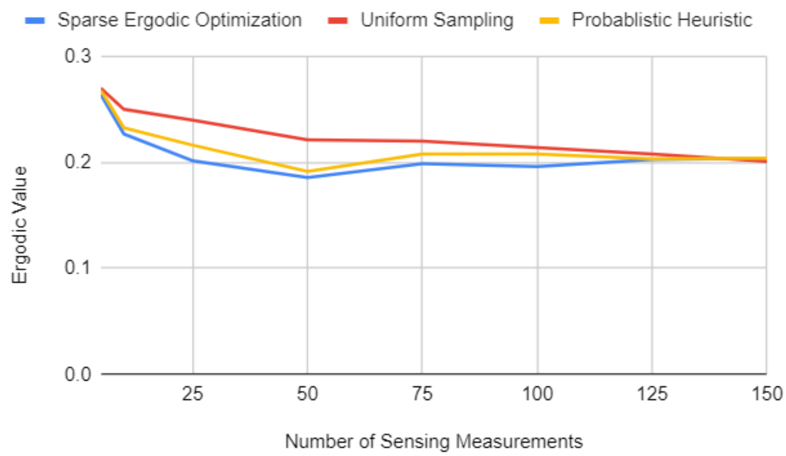

Figure 5.13: **Multi-Agent Experimental Results on Slope Maps:** Results show that sparse ergodic optimization has better coverage performance in terms of the ergodic metric when compared to standard ergodic optimization with uniformly distributed sparse measurements. The probabilistic heuristic results in comparable performance.

5.3 Mixed-Integer Program Relaxation

As described in Section 3.1, the sparse ergodic optimization problem Eq. 3.6 is a mixed integer programming problem. Since mixed integer programming problems are expensive to solve, we relax this formulation by defining $\lambda(t)$ to be a bounded continuous variable $\lambda(t) \in [0, 1]$, and map the resultant values to discrete values in $\{0, 1\}$. We use a solver to compare the results of the mixed integer programming formulation to our relaxed problem to show that this relaxation greatly improves the computational cost of the optimization, without negatively impacting performance.

We compare the performance statistics and computation cost of our approach to that of solving the mixed integer programming formulation of the sparse ergodic optimization problem. We see in Fig 5.14 that our relaxation of the sparse ergodic optimization problems leads to much lower computational costs (i.e. lower run times), while retaining comparable coverage performance in terms of the ergodic metric.

In the multi-agent case, directly solving the sparse ergodic optimization problem as a a mixed integer programming problem becomes intractable as the number of agents increases. As the optimization problem scales, computational costs inherently increase.

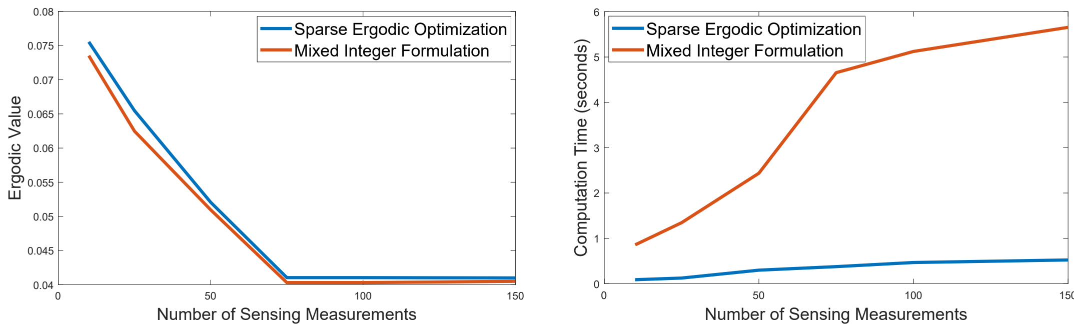

Figure 5.14: **Comparison of Mixed Integer Optimization Problem and Relaxed Sparse Optimization Problem:** Solving the mixed integer formulation of the sparse sensing problem leads to slightly better coverage performance (a), in terms of the ergodic metric, but has a much higher computational cost (b).

This increase in both required computation and memory can grow exponentially with increase in the number of integer variables. Since the multi-agent case introduces T new integer variables for each agent, where T is the length of the trajectory, there is a large increase in the number of additional constraints on the optimization problem for each additional agent. As a result, employing a continuous relaxation as described above makes planning sparse ergodic trajectories for multi-agent teams computationally feasible, while maintaining coverage performance.

Chapter 6

Conclusions

In this thesis we investigate the idea that simultaneously reasoning about where a robot should go and when and where it should take sensing or sampling measurements improves the quality of measurements taken, and the coverage performance achieved by the robot. To this end, we formulate a novel approach to sensing-resource limited coverage by extending the ergodic optimization problem to jointly optimize for both the sensing trajectory and the decision of where to take sensing measurements. Specifically, the set of sensor measurements is posed as a sparse ergodic optimization problem, where the choice to take a measurement is encoded in a vector of sample weights.

We demonstrate the efficacy of sparse ergodic trajectories through experimental evaluation for single and multi-agent coverage of both synthetically generated Gaussian information maps, and information maps built from real-world data. We observe that the joint optimization of trajectory and sensing measurement or sampling locations improves coverage performance in these experimental settings, and infer that this is a function of improvements in informativeness of the measurements taken, and a reduction in overall control cost of the system caused by the joint optimization formulation. Our set of experiments show that there exists a required set of sensing measurements in order to fully characterize a given coverage scenario, in both single and multi-agent cases. The number of required measurements is different for different scenarios. Finally, we infer that there exists a cross-effect between an agent's trajectory and sensing decisions, which make it important to jointly optimize these in cases with limited sensing measurements. This effect is stronger for multi-agent scenarios, leading to larger performance improvements for multi-agent sparse ergodic optimization.

We assume static information maps in this work, that is, we assume that the information distribution over the region being covered stays the same throughout the trajectory's execution. However, in many cases information maps are dynamic and can change over time or with new information from sensor measurements. Future work will look at adapting sparse ergodic optimization for these dynamic settings. This work also assumes the availability of accurate *a priori* information maps, which is not the case for many real-world coverage applications. Future work will seek to use sparse ergodic optimization in order to identify and account for inaccurate information priors with minimum sensor measurements, and take into account sensor noise. Finally, future work will extend sparse ergodic optimization to account for multiple sensors which read in information from multiple information maps.

Bibliography

[1] Agricultural uav (drones). URL https://www.electrocraft.com/motors-for/agriculture/uav-drones/. (document), 1.1

[2] Imaging spectroscopy and the airborne visible/infrared imaging spectrometer (aviris). *Remote Sensing of Environment*, 65(3):227–248, 1998. ISSN 0034-4257. 4.1.2

[3] Vitaly Ablavsky and Magnus Snorrason. Optimal search for a moving target - A geometric approach. In *AIAA Guidance, Navigation, and Control Conference and Exhibit*. AIAA, aug 2000. 2.1

[4] Ian Abraham, Anastasia Mavrommati, and Todd Murphey. Data-Driven Measurement Models for Active Localization in Sparse Environments. In *Robotics: Science and Systems XIV*. Robotics: Science and Systems Foundation, June 2018. doi: 10.15607/RSS.2018.XIV.045. 2.3

[5] Ian Abraham, Ahalya Prabhakar, and Todd D. Murphey. An Ergodic Measure for Active Learning From Equilibrium. *IEEE Transactions on Automation Science and Engineering*, 18(3):917–931, July 2021. ISSN 1558-3783. doi: 10.1109/TASE.2020.3043636. 2.3

[6] Elif Ayvali, Alexander Ansari, Long Wang, Nabil Simaan, and Howie Choset. Utility-guided palpation for locating tissue abnormalities. *Robotics and Automation Letters*, 2017. 2.1

[7] Elif Ayvali, Hadi Salman, and Howie Choset. Ergodic coverage in constrained environments using stochastic trajectory optimization. In *International Conference on Intelligent Robots and Systems*, pages 5204–5210. IEEE, 2017. 2.1, 2.3

[8] Joseph L Baxter, EK Burke, Jonathan M Garibaldi, and Mark Norman. Multi-robot search and rescue: A potential field based approach. In *Autonomous robots and agents*, pages 9–16. Springer, 2007. 2.1

[9] Andreas Bircher, Mina Kamel, Kostas Alexis, Helen Oleynikova, and Roland Siegwart. Receding horizon "next-best-view" planner for 3d exploration. page 1462–1468. IEEE Press, 2016. doi: 10.1109/ICRA.2016.7487281. URL https://doi.org/10.1109/ICRA.2016.7487281. 2.2

[10] Paul S. Blaer and Peter K. Allen. Data acquisition and view planning for 3-d modeling tasks. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 417–422, 2007. doi: 10.1109/IROS.2007.4399581. 2.2

[11] Alberto Candela Garza. *Bayesian models for science-driven robotic exploration*. PhD thesis, 2021. 4.1.2, 4.1.2

[12] Howie Choset. Coverage for robotics–a survey of recent results. *Annals of mathematics and artificial intelligence*, 31(1):113–126, 2001. 2.1

[13] H. Fujisada, F. Sakuma, A. Ono, and M. Kudoh. Design and preflight performance of aster instrument protoflight model. *IEEE Transactions on Geoscience and Remote Sensing*, 36(4), 1998. 4.1.2

[14] L. Hamlin, R. O. Green, P. Mouroulis, M. Eastwood, D. Wilson, M. Dudik, and C. Paine. Imaging spectrometer science measurements for terrestrial ecology: Aviris and new developments. In *2011 Aerospace Conference*, 2011. 4.1.2

[15] Marin Kobilarov. Cross-entropy motion planning. *The Int. J. Robot. Res.*, 31 (7):855–871, 2012. 2.3

[16] Basaran Bahadir Kocer, Boon Ho, Xuanhao Zhu, Peter Zheng, André Farinha, Feng Xiao, Brett Stephens, Fabian Wiesemüller, Lachlan Orr, and Mirko Kovac. Forest drones for environmental sensing and nature conservation. In *2021 Aerial Robotic Systems Physically Interacting with the Environment (AIRPHARO)*, pages 1–8, 2021. doi: 10.1109/AIRPHARO52252.2021.9571033. 1.1

[17] Pablo Lanillos, Seng Keat Gan, Eva Besada-Portas, Gonzalo Pajares, and Salah Sukkarieh. Multi-uav target search using decentralized gradient-based negotiation with expected observation. *Information Sciences*, 282:92–110, 2014. 2.1

[18] Fangchang Ma, Luca Carlone, Ulas Ayaz, and Sertac Karaman. Sparse sensing for resource-constrained depth reconstruction. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 96–103, 2016. doi: 10.1109/IROS.2016.7759040. 2.2

[19] Leonardo Marín, Marina Vallés, Ángel Soriano, Ángel Valera, and Pedro Albertos. Multi sensor fusion framework for indoor-outdoor localization of limited resource mobile robots. *Sensors*, 13(10):14133–14160, 2013. ISSN 1424-8220. doi: 10.3390/s131014133. URL `https://www.mdpi.com/1424-8220/13/10/14133`. 2.2

[20] George Mathew and Igor Mezić. Metrics for ergodicity and design of ergodic dynamics for multi-agent systems. *Physica D: Nonlinear Phenomena*, 240(4): 432–442, 2011. 2.1, 2.3, 2.3

[21] Lauren M Miller and Todd D Murphey. Trajectory optimization for continuous ergodic exploration. In *American Control Conf. (ACC), 2013*, pages 4196–4201. IEEE, 2013. 2.3, 2.3

[22] Lauren M Miller, Yonatan Silverman, Malcolm A MacIver, and Todd D Murphey. Ergodic exploration of distributed information. *IEEE Transactions on Robotics*, 32(1):36–52, 2015. 2.3

[23] Lauren M Miller, Yonatan Silverman, Malcolm A MacIver, and Todd D Murphey. Ergodic exploration of distributed information. *IEEE Transactions on Robotics*, 32(1):36–52, 2016. 2.1

[24] Arturo Rankin, Mark Maimone, Jeffrey Biesiadecki, Nikunj Patel, Dan Levine, and Olivier Toupet. Driving curiosity: Mars rover mobility trends during the first seven years. In *2020 IEEE Aerospace Conference*, pages 1–19, 2020. 1.1

[25] Mark Schmidt, Alexandru Niculescu-Mizil, Kevin Murphy, et al. Learning graphical model structure using l1-regularization paths. In *AAAI*, volume 7, pages 1278–1283, 2007. 3.1

[26] Shrijana Vaidya, Marten Schmidt, Peter Rakowski, Norbert Bonk, Gernot Verch, Jürgen Augustin, Michael Sommer, and Mathias Hoffmann. A novel robotic chamber system allowing to accurately and precisely determining spatio-temporal co2 flux dynamics of heterogeneous croplands. *Agricultural and Forest Meteorology*, 296:108206, 2021. 1.1

[27] Laurence A Wolsey. Mixed integer programming. *Wiley Encyclopedia of Computer Science and Engineering*, pages 1–10, 2007. 3.3

[28] El-Mane Wong, Frédéric Bourgault, and Tomonari Furukawa. Multi-vehicle bayesian search for multiple lost targets. In *International Conference on Robotics and Automation*, pages 3169–3174. IEEE, 2005. 2.1

[29] Gin-Der Wu, Zhen-Wei Zhu, and Chung-Wei Chien. Sparse-sensing-based wall-following control design for a mobile-robot. In *2016 IEEE International Conference on Control and Robotics Engineering (ICCRE)*, pages 1–5, 2016. doi: 10.1109/ICCRE.2016.7476144. 2.2