Sparse Sensing in Ergodic Optimization

Ananya Rao[1], Ian Abraham[2], Guillaume Sartoretti[3], and Howie Choset[1]

[1] Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15203, USA
WWW home page: `biorobotics.org`
[2] Department of Mechanical Engineering, Yale University, New Haven, CT, 06520–8292, USA.
[3] Mechanical Engineering Dept., National University of Singapore, SG 117575.

Abstract. This paper presents a novel, sparse sensing motion planning algorithm for autonomous mobile robots in resource limited coverage problems. Optimizing usage of limited resources while effectively exploring an area is vital in scenarios where sensing is expensive, has adverse effects, or is exhaustive. We approach this problem using ergodic search techniques, which optimize how long a robot spends in a region based on the likelihood of obtaining informative measurements which guarantee coverage of a space. We recast the ergodic search problem to take into account when to take sensing measurements. This amounts to a mixed-integer program that optimizes *when* and *where* a sensor measurement should be taken while optimizing the agent's paths for coverage. Using a continuous relaxation, we show that our formulation performs comparably to dense sampling methods, collecting information-rich measurements while adhering to limited sensing measurements. Multi-agent examples demonstrate the capability of our approach to automatically distribute sensor resources across the team. Further comparisons show comparable performance with the continuous relaxation of the mixed-integer program while reducing computational resources.

Keywords: Search and coverage, Control, Ergodic search

1 Introduction

There is a lack of capability of current robotic systems to explore and search for information in resource limited settings. Such a problem occurs in several applications including planetary science robots with expensive sensors, agricultural robots with limited available measurements, etc. Currently, arbitrary heuristics are used to prioritize resources, but this approach does not consider the dependence on what information the robot has acquired and its mobility capabilities. There is a need for methods that can jointly optimize how robots search and when to take measurements.

Prior works in coverage, especially those based on ergodic search processes, assume that the sensor being used is consistently taking measurements, and therefore, constantly uses energy. In robotic applications, taking a sensor measurement incurs a cost, which this work strives to minimize without compromising the efficacy of autonomous coverage, in terms of the ergodic metric. To improve the applicability and efficacy of robots, especially in resource limited environments, we need to consider optimizing over when to record and acquire sensor measurements (for example, in Fig 1). In this paper, we pose the problem of acquiring a set of optimal 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.

Much prior work in resource-limited robotics focuses on using sparse data to achieve goals like localization [1] and depth reconstruction [2]. However, these formulations pose complex optimizations, which are hard to solve or require forming heuristics that neglect optimality of sensor measurement placement. Unlike many other information-based search methods, trajectories planned using ergodic search processes can be optimized over longer time horizons, and drive the robot to spend time in areas of the search domain in proportion to the amount of expected information at those locations, thus efficiently balancing exploration and exploitation. Optimizing the ergodic metric ties to better coverage of the *a priori* information distribution. Furthermore, ergodic search processes yield special structure in the formulation of the optimization problem, which we leverage to formulate sampling decision times that are solutions to a convex optimization problem. The sampling solutions result in a near-optimal set of sensing measurements along the robot's trajectory.

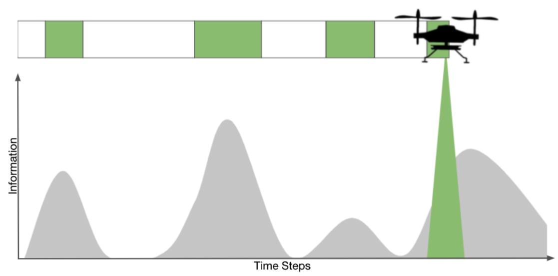

Fig. 1: Proposed Approach for Sparse Ergodic Optimization: Our approach automates how a robot explores an area and what informative measurements to collect in a joint optimization problem. Illustrated above is an example sparse sensing solution for covering a one-dimensional information distribution (gray distribution), where peaks correspond to areas of high expected information, and the green colored bar represents when the robot takes a measurement.

The organization of this paper is as follows: In Section 2 we discuss existing search and coverage methods, as well as recent advances in sparse sensing. A brief background on ergodic search processes is also presented in Section 2. Details on our sparse ergodic optimization approach, and its application to distributing sensing resources across a multi-agent team are discussed in Section 3. In Section 4 we present and discuss the results of our systematic set of experiments, in which we observe that our approach can maintain coverage performance while significantly decreasing sensing costs, in both single and multi-agent settings. Finally, we presenting concluding remarks in Section 5.

2 Prior Work

Search and Coverage Methods: Search and coverage methods have widespread uses in a number of important applications, including search and rescue, cleaning robots, and automated agriculture techniques. Both object search and coverage involve motion planning, typically based on an *a priori* information distribution in a domain. Most search and coverage techniques involve the construction and use of some kind of information map which encodes the utility of the robot exploring specific regions of the search space.

Broadly, area search and coverage methods fall into three categories: geometric, gradient-based and trajectory optimization-based approaches. Geometric methods, like lawnmower search, are good techniques for uniformly covering a search domain [3, 4]. Such approaches are useful when no *a priori* information is provided, or when the *a priori* information is presumed to be a uniform probability distribution over the search domain. On the other hand, since these approaches do not utilize any *a priori* information about regions of interest, the efficacy of exploration and object search is upper bounded.

When *a priori* information is available, a representation of how information is distributed across a search domain (i.e. an *a priori* information map or information distribution) can be leveraged to develop better search and coverage techniques. In gradient-based (i.e. "information surfing") methods [5, 6, 7], the short-term information gain is greedily maximized by guiding the search agent in the direction of the derivative of the information map around its position. Since the agent is always driven in the direction of greatest information gain, it is led to areas with higher likelihood of finding objects. However, these approaches can leave areas unexplored, because they generally do not take into account the uncertainty associated with the *a priori* information distribution, which can help differentiate between unexplored areas with low information and areas with no information to be gained. These approaches are very sensitive to noise in the information map. This is because noise leads to inaccurate gradient estimations, which in turn leads to agents greedily over-exploiting local information maxima.

Trajectory optimization-based methods formulate search as an information gathering maximization problem, which is then solved by planning a path for the search agent [8]. The cost function that drives the optimization in many of these approaches combines both the information distribution and its associated uncertainty, thereby reducing the likelihood of leaving areas unexplored.

Sparse Sensing: 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 [1], depth reconstruction [2] and wall-following [9]. 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 decided where to take these limited measurements.

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 [10, 11]. 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 [12, 11]. Ergodic search is able to balance both exploration and exploitation goals by accounting for both diffused information densities and highly focused points of information [13, 8, 14]. 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 [8] produces trajectories for agents, such that they spend time in each area 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 1 and the underlying information map.

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

where γ is the agent's trajectory, defined as $\gamma : (0, t] \to \mathcal{X}$, t is the discrete time horizon, and δ is the Dirac delta function, with $\mathcal{X} \subset \mathrm{I\!R}^{\mathrm{d}}$ 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 [8]:

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

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 the 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 $f : \mathcal{Q} \times \mathcal{U} \to \mathcal{TQ}$, such that

$$\begin{aligned} \mathbf{u}^*(t) &= \arg\min_{\mathbf{u}} \Phi(\gamma(t)), \\ \text{subject to } \dot{\mathbf{q}} &= f(\mathbf{q}(t), \mathbf{u}(t)), \quad \|\mathbf{u}(t)\| \leq u_{max} \end{aligned} \tag{3}$$

where $\mathbf{q} \in \mathcal{Q}$ is the state and $\mathbf{u} \in \mathcal{U}$ denotes the set of controls. Eq.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 [14], 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.

3 Sparse Sensing in Ergodic Optimization

We approach sparse sensing in ergodic coverage in two steps. First, we extend ergodic optimization to formulate a sparse ergodic optimization problem to jointly optimize for both an agent's trajectory, and the importance of taking a sensor measurement at each point along the trajectory. Second, we apply this formulation to ergodic coverage using multi-agent teams, and jointly optimize the trajectories and sensing decisions of each agent in the team, resulting in the optimal distribution of limited sensing resources amongst agents.

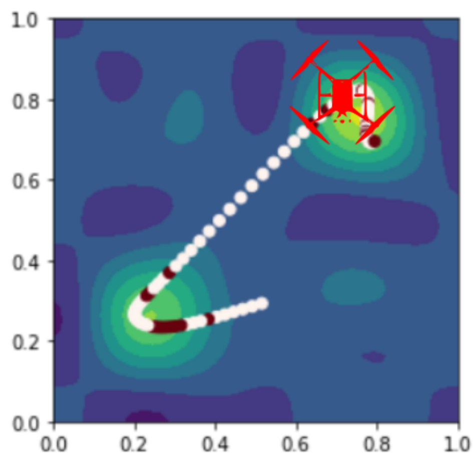

Fig. 2: Sparse Ergodic Trajectory Example: Trajectories are generated with a decision variable for when to take a sensor measurement. This approach automatically updates trajectories which provide maximal information gathering with minimal sensing resources.

3.1 Sparse Ergodic Optimization

In many search and coverage applications, only a few sensor measurements or data points are required to fully comprehend and optimize a model. Note that in the optimization in Eq 3, the locations at which sensor measurements are obtained are uniform in time, that is, each measurement along the trajectory is treated as being equally important and contributing equally to minimizing the ergodic metric. In many scenarios, due to both the non-uniform distribution of information in the region being covered, as well as the dynamic constraints of the robotic agents being used, this can lead to extraneous measurements that do not contribute to reconstructing the information prior.

For example, if there are few regions of high information, which are separated by larger regions of low expected information, uniform sensor measurements would result in many measurements that provide no information gain (see Fig. 2). In this case, ergodic performance would be improved by only taking sensor measurements near the areas of higher information.

In this work we extend ergodic optimization in order to find this set of optimal measurements for a given search scenario. We optimize for the set of optimal sensor measurements by posing the following ergodic optimization,

$$\begin{aligned} \mathbf{u}^*(t), \lambda^*(t) &= \arg\min_{\mathbf{u},\lambda} \Phi'(\gamma(t)), \\ \text{subject to } \dot{\mathbf{q}} &= f(\mathbf{q}(t), \mathbf{u}(t)), \quad \|\mathbf{u}(t)\| \leq u_{max} \end{aligned} \tag{4}$$

where $\mathbf{q} \in \mathcal{Q}$ is the state, $\mathbf{u} \in \mathcal{U}$ denotes the set of controls, and $\lambda(t) \in \{0, 1\}$. Note that here we define λ as a decision variable. We promote sparsity in the sample measurements by regularizing λ with an L^1 optimization which promotes sparsity [16].

We augment the ergodic metric in Eq. 2 in the following manner

$$\Phi'(\gamma(t)) = \sum_{k=0}^{m} \alpha_k \left| c_k(\gamma(t), \lambda(t)) - \xi_k \right|^2 + \sum |\lambda_k|, \tag{5}$$

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 spatial time-average statistics of the agent's trajectory (from Eq. 1) are also modified to be,

$$C'^t(\mathbf{x}, \gamma(t)) = \frac{1}{\sum_t \lambda(t)} \sum_0^t \lambda(t) \delta(\mathbf{x} - \gamma_i(\tau)), \tag{6}$$

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

$\lambda(t)$ represents the decision variable for choosing whether to take a sensor measurement or not at a given location in the search domain. Thus far, $\lambda(t)$ is defined to be an integer (i.e. $\lambda(t) \in \{0, 1\}$), resulting in Eq 4 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 [17], and due to requiring direct search methods that do not scale with longer time horizons. For this reason, we employ a relaxation of the problem Eq 4 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 4, and then map the resultant continuous values of $\lambda(t)$ to discrete values $\{0, 1\}$. The procedure for optimizing the sparse ergodic problem Eq 4 is depicted in Fig 3.

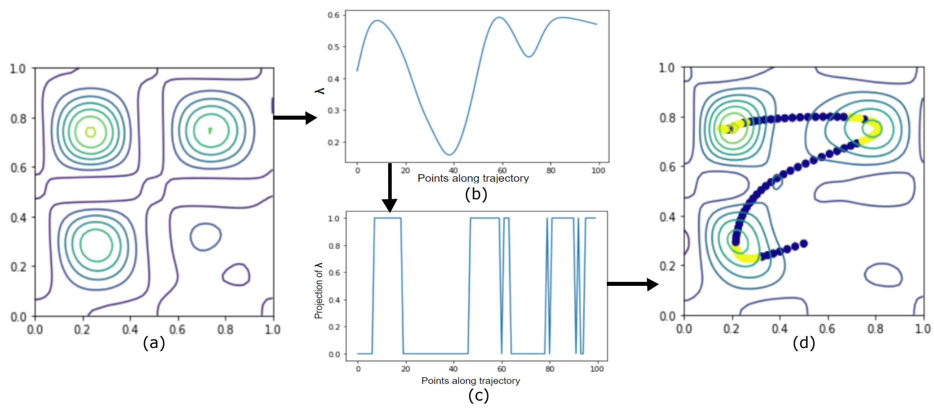

Fig. 3: Mixed-Integer Continuous Relaxation for Sparse Ergodic Optimization: Illustrated is the process with which we jointly optimize for trajectories and when to sample. The decision variable λ is relaxed to be continuous $[0, 1]$ where solutions are projected into the integer $0, 1$ space. As input, 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.

When we investigate the numerical values of the decision variables $\lambda(t)$ calculated in sparse ergodic optimization Eq. 4, we see that there are peaks formed that correspond to peaks in information in the *a priori* information distribution (Fig 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.

3.2 Multi-Agent Sparse Ergodic Optimization

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. We apply sparse ergodic optimization (Eq. 4), as described in Section 3.1 to distribute sensing resources amongst a multi-agent team. A sample result is shown in Fig 4.

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

$$C'^t(\mathbf{x}, \gamma(t)) = \frac{1}{Nt\sum_{i=1}^{N}\sum_t \lambda_i(t)} \sum_0^t \lambda_i(t)\delta(\mathbf{x} - \gamma_i(\tau)), \tag{7}$$

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

We use the augmented ergodic metric defined in Eq 5 to drive the joint optimization of the set of agent trajectories and sensing decision variables. The ergodic optimization posed in Eq 4 is used to generate optimal controls and set of optimal sensor measurements for each agent.

We again employ a relaxation of the problem described by Eq 4 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.

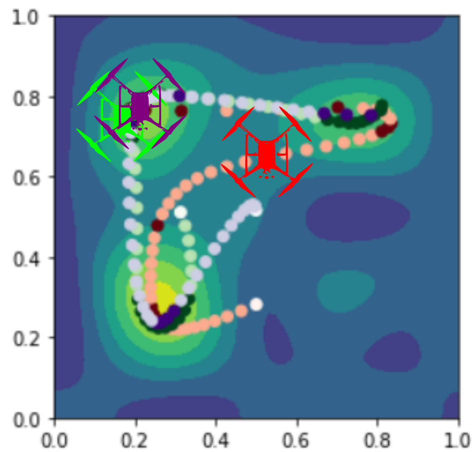

Fig. 4: Mutli-Agent Sparse Ergodic Trajectory Example: Trajectories for a multi-agent team are generated with a decision variable for when to take a sensor measurement. This approach automatically updates trajectories which provide maximal information gathering with minimal sensing resources.

4 Results and Discussion

In this section we empirically show that for a specific ergodic coverage problem, there exists a minimal set of sensor measurements to take in order to gain enough information to fully characterize and solve the coverage task. Through our experiments we illustrate that optimizing for this set of samples, or sensor measurements, leads to improvements in ergodicity for coverage problems. This supports our intuition that in many coverage scenarios, extraneous samples are taken that either do not contribute to or negatively impact the optimization. Further, we show that our sparse ergodic optimization approach can be applied to multi-agent coverage tasks in order to optimally distribute a given sensing budget amongst all of the agents. In multi-agent coverage scenarios as well, we empirically show that for a specific ergodic coverage problem, there is a minimal set of sensor measurements required to cover the information prior.

4.1 Experimental Details

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

$$x(t+1) = f(x,u) = x(t) + \tanh(u) \tag{8}$$

where $x(t) \in \mathbb{R}^2$ and $u \in \mathbb{R}^2$ is constrained by the tanh function to be 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

$$p(x) = \sum_{i=1}^{3} \eta_i \exp(||x - c_i||^2_{\sum_i^{-1}}) \tag{9}$$

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.

As described in Section 3.1, the sparse ergodic optimization problem Eq. 4 is a mixed integer programming problem. Since mixed integer programming problems are difficult 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 investigate the performance of the sparse ergodic optimization described in Section 3.1 in terms of coverage performance, using ergodicity as the metric. We compare our coverage results to two different approaches. The first is standard ergodic optimization Eq. 2, where sensor measurements are uniformly distributed along the robot's trajectory. The second 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.

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.

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. Agents starting positions, initial information maps, and sensing budgets are kept identical among experiments with different controllers to ensure that our results are comparable.

4.2 Single Agent Sensing Distribution

When looking at results of the sparse ergodic optimization approach in terms of overall coverage performance, measured through the ergodic metric, we observe that there is a minimal number of sensor measurements to be taken to minimize ergodicity (see Fig 6). 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.

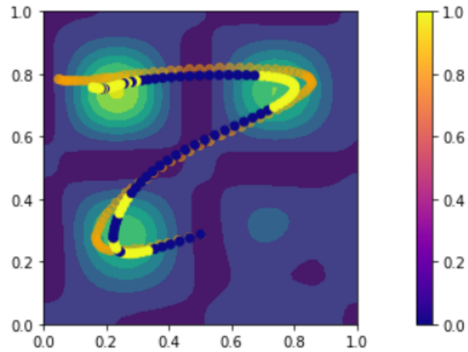

Fig. 5: Standard Ergodic and Sparse Ergodic Trajectory Comparison: Standard ergodic optimization results in a different final trajectory (orange trajectory) than jointly optimizing for both trajectory and sensor weight (blue and yellow trajectory). This implies a cross-effect between agent position and if a sensing measurement is taken, due to which joint optimization becomes more important as the coverage problem scales.

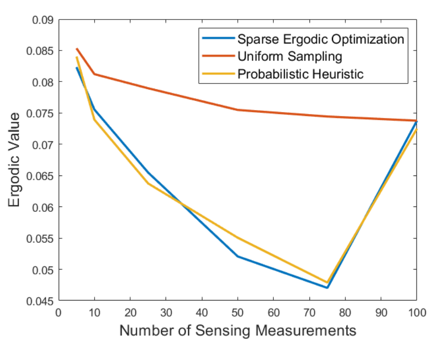

Fig. 6: Comparison of Baseline Approaches and Sparse Optimization: Results show that sparse ergodic optimization has better coverage performance in terms of the ergodic metric when compared to standard ergodic optimization with uniformly distributed sparse measurements. Using the probabilistic heuristic results in comparable performance.

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, the probabilistic heuristic 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), 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.

Further, 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 7 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.

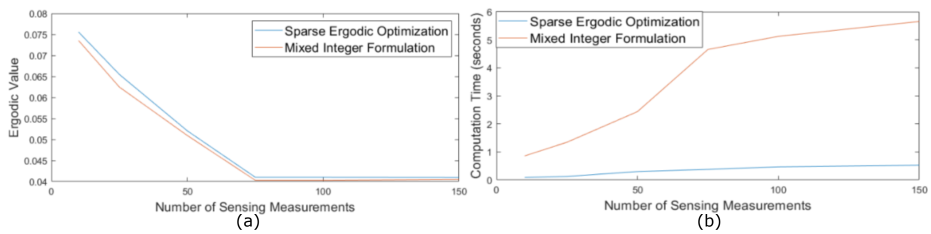

Fig. 7: Comparison of Mixed Integer Optimization Problem and Relaxed Sparse Optimization Problem: Solving the mixed integer formulation of the sparse sensing problem leads to slightly better coverage performance (a), in terms of the ergodic metric, but has a much higher computational cost (b).

4.3 Multi-Agent Sparse Ergodic Optimization

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.

Similar to our single agent sparse ergodic optimization results, we see than 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.

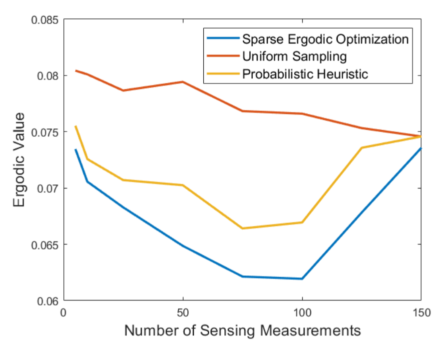

Fig. 8: Comparison of Baseline Approaches and Sparse Optimization for Multi-Agent Coverage: Results show that multi-agent sparse ergodic optimization has better coverage performance in terms of the ergodic metric when compared to standard ergodic optimization with continuous sensing, and with uniformly distributed sparse measurements.

for lower sensing budgets) (see Fig 8). 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 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 8). As explained in in Section 4.2, 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.

5 Conclusion

In this paper we investigate the idea that there is an optimal set of sensing measurements that can be taken during coverage, in order to fully characterize the problem, reduce sensing costs and avoid local sensing optima. To this end, we formulate a novel approach to sensing-resource limited coverage by modifying ergodic optimization to jointly optimize for both the sensing trajectory and the decision of where to take sensing measurements. Specifically, the set of optimal 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. Our set of experiments show that there exists an optimal set of sensing measurements for a given coverage scenario, in both single and multi agent cases, using ergodicity, a state-of-the-art coverage metric. We also 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.

This work experimentally shows that there are improvements to be made in resource-limited coverage scenarios through sparse sensing. Specifically, we show that there exists an optimal set of measurements required to solve a coverage problem. Future work will focus on proving theoretical guarantees regarding the existence of this optimal set of sensing measurements. This work 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.

Bibliography

[1] Marín, L., Vallés, M., Soriano, A., Valera, A., Albertos, P.: Multi sensor fusion framework for indoor-outdoor localization of limited resource mobile robots. Sensors **13**(10), 14,133–14,160 (2013). DOI 10.3390/s131014133. URL https://www.mdpi.com/1424-8220/13/10/14133

[2] Ma, F., Carlone, L., Ayaz, U., Karaman, S.: Sparse sensing for resource-constrained depth reconstruction. In: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 96–103 (2016). DOI 10.1109/IROS.2016.7759040

[3] Ablavsky, V., Snorrason, M.: Optimal search for a moving target - A geometric approach. In: AIAA Guidance, Navigation, and Control Conference and Exhibit. AIAA (2000)

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

[5] Lanillos, P., Gan, S.K., Besada-Portas, E., Pajares, G., Sukkarieh, S.: Multi-uav target search using decentralized gradient-based negotiation with expected observation. Information Sciences **282**, 92–110 (2014)

[6] Baxter, J.L., Burke, E., Garibaldi, J.M., Norman, M.: Multi-robot search and rescue: A potential field based approach. In: Autonomous robots and agents, pp. 9–16. Springer (2007)

[7] Wong, E.M., Bourgault, F., Furukawa, T.: Multi-vehicle bayesian search for multiple lost targets. In: International Conference on Robotics and Automation, pp. 3169–3174. IEEE (2005)

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

[9] Wu, G.D., Zhu, Z.W., Chien, C.W.: Sparse-sensing-based wall-following control design for a mobile-robot. In: 2016 IEEE International Conference on Control and Robotics Engineering (ICCRE), pp. 1–5 (2016). DOI 10.1109/ICCRE.2016.7476144

[10] Miller, L.M., Silverman, Y., MacIver, M.A., Murphey, T.D.: Ergodic exploration of distributed information. IEEE Transactions on Robotics **32**(1), 36–52 (2015)

[11] Abraham, I., Mavrommati, A., Murphey, T.: Data-Driven Measurement Models for Active Localization in Sparse Environments. In: Robotics: Science and Systems XIV. Robotics: Science and Systems Foundation (2018). DOI 10.15607/RSS.2018.XIV.045

[12] Abraham, I., Prabhakar, A., Murphey, T.D.: An Ergodic Measure for Active Learning From Equilibrium. IEEE Transactions on Automation Science and Engineering **18**(3), 917–931 (2021). DOI 10.1109/TASE.2020.3043636

[13] Ayvali, E., Salman, H., Choset, H.: Ergodic coverage in constrained environments using stochastic trajectory optimization. In: International Conference on Intelligent Robots and Systems, pp. 5204–5210. IEEE (2017)

[14] Miller, L.M., Murphey, T.D.: Trajectory optimization for continuous ergodic exploration. In: American Control Conf. (ACC), 2013, pp. 4196–4201. IEEE (2013)
[15] Kobilarov, M.: Cross-entropy motion planning. The Int. J. Robot. Res. **31**(7), 855–871 (2012)
[16] Schmidt, M., Niculescu-Mizil, A., Murphy, K., et al.: Learning graphical model structure using l1-regularization paths. In: AAAI, vol. 7, pp. 1278–1283 (2007)
[17] Wolsey, L.A.: Mixed integer programming. Wiley Encyclopedia of Computer Science and Engineering pp. 1–10 (2007)