

# A Bounded Suboptimal Environmental Monitoring Algorithm

Elizabeth A. Ricci<sup>1</sup> and Ross A. Knepper<sup>1</sup>

**Abstract**—Autonomous robot monitoring tasks require the robot to traverse a map over an extended period of time for the purpose of monitoring some value of interest. The robot must choose paths that efficiently move through the environment and perform observations, such that the maximum amount of information is gained. The robot uses these observations to predict the state of the entire environment, including unobserved portions, using a low-rank model built from prior recorded data. We introduce a path planner that selects a set of points in the map that maximize the entropy subject to a cost constraint. We evaluate our method in simulation on the problem of predicting fall foliage colors using a satellite image dataset and a UAV traversal of the planned path.

## I. INTRODUCTION

Most robots collect information about their environment but are constrained to sense only a limited area around the location of the robot. In many problems, like environmental monitoring [8, 13, 14, 17] or persistent monitoring [17, 18, 21], the mission involves sensing broadly in order to build up a map of some quantity throughout a larger region. Consequently, the robot must navigate a path to maximize information gain. In the absence of any structure in the data, building a map amounts to a coverage problem. Often though, there is structure in the data that can be exploited to at least estimate a complete map from limited observations. In this paper, we seek to plan a path for a robot through a region that maximizes the information gathered while adhering to cost constraints.

The measured quantity can have structural correlations over space and time, such as recurring patterns over separate regions or cyclic seasonal repetition. Example measured quantities include vegetation health [20], ocean properties [13, 14, 17, 17], hurricane predictors [8], survivor locations [16], and wireless signal strength [4]. A key insight of this work is that not all areas of a region hold as much uncertainty as others, thus providing less information. Consequently, most of the map may be reconstructed by observing a minority of the locations within it. To make initial decisions about where to visit, we consider the case that there is some prior data about the region to be monitored that enables entropy computations. Our method maximizes information gain by sampling points that have the greatest joint entropy conditioned on prior observations.

### A. Problem Statement

A robot travels around a map comprising a finite set of  $p$  locations that constitute possible sampling points. Each

point in the set holds a scalar-valued, time-varying parameter. At the time the robot makes each observation, it learns the true value of the parameter at its location at that time. We assume the time-varying rate of the quantity is negligible compared with the time elapsed in traveling around the map to make observations. A cost function encodes any desired path qualities, such as end point, and travel cost from one location to another.

In addition to the map structure, the robot is also given some prior map data from  $t$  earlier times. The prior data consists of snapshots of the map (or subsets thereof) each at a single point in time. From these observations, one can estimate a joint probability density function for any subset of the random variables corresponding to those observations.

The *offline problem* is to plan a path through a subset of locations that maximizes expected information gain overall. The online planning problem is not addressed in this paper.

In order to illustrate the techniques of our contribution, we focus on using prior satellite data of eastern North America to predict fall foliage colors. Some areas in each prior map are occluded by cloud cover. We model the robot as a UAV that travels around the region captured by the map and makes observations about the color of observed pixels in the map. The cost of a traversal is path length, measured as the sum of Euclidean distances between consecutive observations.

### B. Related Work

Information-based path planning problems concern the maximization of some information-related function subject to some constraints. The specific functions and constraints are selected based on the desired application. Our work focuses on a relatively simple problem: the maximization of information gain subject to path length constraints. This problem is similar to the problems studied in [4, 9, 16, 21]. Our chosen problem formulation lends itself to high-level solutions where specific motion decisions are deferred to the controller. As a result, the methods used are fairly general and can be used on a wide variety of robotic systems. Some approaches explicitly incorporate real-world constraints, such as obstacles, environmental conditions and robot kinematics into the planning problem, passing more control to the planning algorithm [1, 12, 13, 15, 22]. While our approach results a path containing way-points for a robotic system to sample from, solutions to the aforementioned problems are in the form of a lower-level motion plan. This gain in control allows for more performant solutions for application-specific planners, but requires more initial information and results in the loss of flexibility.

\*This work was supported by the Air Force Office of Scientific Research under award number FA9550-17-1-0109.

<sup>1</sup>The authors are with the Department of Computer Science, Cornell University, Ithaca, NY, USA {ericci, rak}@cs.cornell.edu

In the information-based planning domain, we want to minimize the uncertainty of the state of some object(s). Uncertainty is typically quantified using a form of entropy or information gain. These information functions can be computed using prior information about the state of the object(s) considered. There are two primary ways to model this prior information, which can then be used to compute the information functions. In our approach we assume that our data is normally distributed and represent it using a Gaussian process, this is similar to methods used in [2, 10, 11, 13, 14]. In cases where the data is not normally distributed, Bayesian models can be used to model information containing many variables. This is seen in work where the information function is represented in terms of a non-Gaussian sensor objective [1, 12, 22]. In the sensor based approach, the Bayesian model encodes the probability mass function of sensor readings. As a result, after taking some sensor readings the conditional entropy measures the amount of uncertainty of the value of all other sensor readings. This differs from our approach, which considers the probability density function of the states of the sampling points. With this representation, the conditional entropy represents the amount of uncertainty about the unseen sampling points after a number of sampling points are observed. In our approach we can more clearly consider the mutual information between a set of targets or sampling points, due to our assumption that a target can be completely observed at some sampling location.

There are two primary approaches used to solve information-based path planning problems: greedy and sampling. Greedy solutions are often used in this domain, due to the submodular property of entropy-based information functions [5, 6, 14, 15]. With minor modifications, classical sampling-based planners can be adapted for use in the information planning domain [4, 7, 12, 22]. Our method uses a greedy approach introduced in [23].

### C. Contributions

We introduce an offline planner which uses the submodularity of entropy to choose a set of sampling points that approximately maximizes the joint entropy while remaining within a given cost budget. We present an algorithm for sampling observations subject to a cost constraint on path traversal. We show that the joint entropy of the path returned by our algorithm is at most  $\frac{1}{2}(1 - e^{-1})$  times the optimal solution.

In order to demonstrate the utility of our general method for predicting the entire region using a small number of samples, we propose the use of a low-rank model to predict the values of unobserved sampling points using the observed points.

In the remainder of this paper, we present our method for computing entropy in Section II and our algorithm that finds an optimal information-gain path in Section III. We show how the information gained by traversing this path can be used to estimate the complete map in Section IV. Then we evaluate the algorithm in Section V and conclude in Section VI.

## II. COMPUTING ENTROPY

Our method requires a submodular set function  $f$  that yields differential entropy given a set of observations of Gaussian random variables. For the following methods we assume that the sample space  $S$  is a finite set of random variables and that each random variable has a set of prior values and a Gaussian distribution.  $S$  is drawn from the map locations indexed as  $\{1, \dots, p\}$ .

### A. Joint Differential Entropy

Joint differential entropy (hereafter called joint entropy) is a measure of the uncertainty of a set of continuous random variables and can be computed from a covariance matrix. Let  $\Sigma$  be a covariance matrix of the random variables in the sample space  $S$  and let  $s'$  be a subset of the random variables in  $S$ . A square submatrix of  $\Sigma$  representing the covariance matrix of  $s'$ , called  $\Sigma_{s'}$ , is composed of the rows and columns of  $\Sigma$  corresponding to the random variables in  $s'$ . Then the joint entropy of the samples in  $s'$  is equal to  $\log \det(\Sigma_{s'})$  [6]. Since the covariance matrix is a symmetric positive definite square matrix, the function  $f = \log \det(\Sigma_{s'})$  is a submodular function.

## III. ALGORITHM

In this section we introduce our offline planner, which is based on the generalized cost-benefit algorithm by Zhang and Vorobeychik [23].

### A. Cost function

The cost function will be used to compute the cost of the shortest path between a set of points and can be constructed to enforce any desired path properties. In many cases, such as when a tour is desired, the cost function is an NP-hard problem. In order to make the planner tractable we allow for approximate cost functions. In order to make an optimality guarantee about the planner the approximate cost function should be a  $v(n)$ -approximation where  $n$  is the number of points in the set. A cost function is a  $v(n)$ -approximation when it is at most  $v(n)$  times greater than the true cost. In addition the cost function should be  $\alpha$ -submodular, that is when  $x, A, B$  are minimized and  $A \subset B$  the quotient  $\frac{\text{cost}(A \cup x) - \text{cost}(A)}{\text{cost}(B \cup x) - \text{cost}(B)}$  is equal to  $\alpha$  [23].

### B. Offline Planner

Given the covariance matrix of the sampling points, and a starting point and a budget, the offline planner returns a set of points to sample such that the joint entropy of the points is approximately maximal and that the approximate cost of a path between the points is within the budget.

The algorithm used to select the sampling points is based on a method introduced by Zhang and Vorobeychik [23], which has the following optimality guarantee: the joint entropy of the points selected is within  $\frac{1}{2}(1 - e^{-1})$  times the maximum joint entropy possible within the cost constraints. We modified the original algorithm to improve its efficiency and supplied joint entropy as the submodular function. The cost constraint was expressed in terms of nearest neighbor distances. Our method is shown in Algorithm 1.

---

**Algorithm 1:** Offline Planner

---

**Data:** covariance\_matrix, start\_point, all\_points, budget  
**Result:** sample\_points

```
1 begin
2   sample_points = [start_point]
3   current_entropy = entropy(sample_points,
4     covariance_matrix)
5   current_cost = 0
6   points = [] # only consider points that are close
7     enough to the start
8   for point  $\in$  all_points do
9     if cost(point, sample_points)  $\leq$  budget then
10      points = points  $\cup$  point
11   while points  $\neq$   $\emptyset$  AND current_cost < budget do
12     for point  $\in$  points do
13       cost_change[point] = cost(point,
14         sample_points) - current_cost
15       entropy_change[point] = entropy(point,
16         sample_points, covariance_matrix) -
17         current_entropy
18       # if the cost of sampling a point is ever
19         over budget, remove it from consideration
20       if cost(point, sample_points) + current_cost
21         > budget then
22         points = points  $\setminus$  point
23     # add the point that maximizes the ratio of
24     entropy change to cost change to the list of
25     sample points
26     new_point = max(entropy_change / cost_change)
27     sample_points = sample_points  $\cup$  new_point
28     points = points  $\setminus$  new_point
29     current_cost = cost(sample_points)
30     current_entropy = entropy(sample_points,
31       covariance_matrix)
32   return sample_points
```

---

#### IV. FALL FOLIAGE MAPPING PROBLEM

##### A. Data

We evaluate our methods on a fall foliage dataset, containing 186 satellite images of the northeast region of North America collected from the fall months over the course of three years collected using NASA Worldview.<sup>1</sup> More specifically there is one  $1629 \times 1076$  pixel image for every day between September 15 and November 15 for the years 2014, 2015 and 2016. Over the course of the three months, the hue of portions of the land in the images changes from green to orange, yellow and brown, as a result of the foliage changes. The original satellite images contain large areas of clouds which obfuscate the hue of the land below. As a result, we model pixels containing clouds as missing data. A support vector machine was used to identify pixels

<sup>1</sup>worldview.earthdata.nasa.gov

containing clouds. In addition, due to the length of time it takes to capture a satellite image, many of the images have obvious hue differences due to lighting changes. We negate the effects of these inconsistencies by splitting affected images into two images on the line of hue shift, where one portion of the image was the section of the original image and the remainder is black representing missing data. The splitting line was identified using Hough transform on the edges of the image. The edges were identified using Canny edge detection. After this processing we have 298 processed images, all containing missing data.

For use in the planning methods, the sampling set of this data set are the set of all pixels in an image. Due to space constraints, we only consider about 100 by 100 pixel subsets of the image. With the path lengths we have chosen to evaluate, this does not limit the planners.

##### B. Cost Function

For our particular dataset we approximate the shortest path using the nearest neighbor algorithm, which is a known-greedy approximation for shortest path problems [3]. Starting with a given point, the algorithm adds points to the path based on proximity to the current point. That is, at each timestep the point to be traveled to is the one that is closest to the current point. Given a set of points  $\{v_1, \dots, v_n\}$ , the nearest neighbor algorithm returns a travel ordering of these points, starting at  $v_1$ , the ordering is notated  $\{v'_1, \dots, v'_n\}$ . The cost of the path is:

$$\sum_{i=1}^{n-1} d(v'_i, v'_{i+1}) \quad (1)$$

where  $d(a, b)$  is the Euclidean distance between the points  $a$  and  $b$ . This cost function is a  $v(x)$ -approximation and  $\alpha$ -submodular by [3].

##### C. Generalized Low Rank Model

The planned path returned by the algorithm in Section III is designed to gain as much information as possible about the overall map within a limited budget. Having collected those observations, we need a technique to complete the map through extrapolation. We construct a low-rank model to predict the full map. In order to compute the model we first form a data matrix  $D$  of  $t$  rows corresponding to the time snapshots and  $p$  columns containing the observed values at each location at the specified time. A corresponding mask matrix identifies entries in  $D$  that were not observed. The problem of constructing a low-rank model is to find two matrices  $X$  and  $Y$  that minimize the quadratic loss function,  $(X^T Y - D)^2$ , subject to a non-negative constraint.

We use a Julia implementation of the generalized low rank model of Udell et al. [19] to factor the data matrix while ignoring the missing data generated by in the processing stage. More specifically, the model takes the form of two matrices, the  $X$  matrix has a column for each image and the number of rows is equal to the rank of the model. The  $Y$  matrix has a column for each pixel and the number of rows is again the rank of the model. The columns of the

$Y$  matrix are the principal components of the data set. A set of principal components for a rank three model can be seen in Fig. 2. The product  $X^T Y$  is an approximation of the original data matrix and provides an approximate value for each missing data point.

#### D. Map Completion

Given a vector  $d_{\text{path}}$  containing the values of the observed pixels in an image, the  $Y$  matrix can be used to estimate the values of the pixels not observed. To do so, we take the set of observed pixels situated in a row vector of length  $p$  that includes unobserved pixels. We compute an ‘‘image digest’’ vector  $x_{\text{path}}$  of dimension equal to the model rank by

$$x_{\text{path}} = (Y^T)^+ d_{\text{path}},$$

where  $(Y^T)^+$  represents the pseudoinverse of  $Y^T$ . We then reconstruct a dense image by

$$\tilde{d}_{\text{dense}} = Y^T x_{\text{path}}.$$

#### E. Rank Selection

There is no straightforward method for determining the optimal rank for a given data set, so we considered a number of ranks and choose the one that minimized the reconstruction error for known images. That is, for the approximated data matrix, the product  $XY^T$ , we computed the reconstruction error by summing the difference in the original data and the corresponding values in the approximated data matrix. The graph in Fig. 1 shows the reconstruction error per pixel for six different models computed of ranks 3, 5 and 11. These three ranks were chosen after a preliminary study found that these ranks resulted in the lowest reconstruction error out of models from ranks 1 to 13. From this study we found that one of the rank 5 models minimized the reconstruction error, this model is used in the evaluation section. We note that the function used to compute these models is non-deterministic, as the model is initialized with random matrices.




Fig. 1: A graph of the error per pixel of the computed generalized low rank models for ranks 3, 5, and 11. Each unit on the horizontal axis represents one trial of factorizing the  $D$  matrix with the chosen rank. The trials are sorted from lowest to highest.

## V. EVALUATION

In this section, we detail our preliminary simulated experimental results.

### A. Baselines

In this subsection we introduce three baseline planners that will be compared with our planner.

1) *Random Baseline*: Given a budget of  $n$ , the random baseline planner randomly selects  $n$  points from the sample point set to visit. This baseline has no path constraint, that is the cost to visit all of the sample points is not considered.

2) *Straight-Path Baseline*: The straight-path planner assumes that the sampling points have a location-based ordering. For the fall foliage dataset the sampling points are pixels and are in raster-scan order starting at the top left corner of the image. Given a starting point with ordering index  $s$  and a budget of  $n$ , the straight-path planner selects the points with indices  $s$  to  $s + n - 1$  to visit.

3) *Sampling-based Baseline*: The sampling-based planner is an any-time rapidly exploring tree based planner, called rapidly exploring information gathering tree (RIG-tree) introduced by Hollinger and Sukhatme [4]. The RIG-tree planner builds a tree from a start node through sampling the sampling point space, adding nodes and edges that are within some distance of each other. More precisely, after a sample is taken, the point that are closest to it and within some delta of the nearest existing vertex is found. This point is then used as a center point of a circle of some fixed radius. All vertices within the circle are extended with an edge to a new vertex that is as close to the point as possible, while being at most delta away from the vertex. Since this is an anytime method, it can be run for any amount of time with results improving with diminishing returns as the time increases.

For our evaluation, the sampling-based baseline method will be run with a delta value of 5 and a radius of 10 for 100 iterations.

### B. Memory Limitations

A primary limitation of our method is the memory cost associated with the covariance matrices necessary to compute the joint entropy of a set of points. Given a cost budget of  $n$ , both our offline planner and the sampling-based baseline planner need the covariance matrix associated with all the sampling points within an  $n$  radius of the starting point. That is,  $O(n^2)$  points will be considered, and the covariance matrix of the set of points is  $O(n^4)$ . Thus, for both planners the theoretical upper bound on memory is  $O(n^4)$ . In the following subsection we show that, on our dataset, the performance of our planner reaches a near optimal reconstruction before the memory needs of the planner become intractable.

### C. Results

To evaluate our method we consider both the joint entropy of the samples selected and reconstruction error. We note that for our specific dataset, our planner is tractable for budgets up to 100.






Fig. 2: Principal components of a rank 3 model.

To compare the joint entropy of the sample sets produced by each method we plan paths with a budget of 50 starting from eight different locations and the planning space is limited to a 100 by 100 pixel subregion. For the random planner, no starting point is given, and 50 points are randomly selected from the full image. The average joint entropy values are shown in Table I.

TABLE I: Joint entropy results for paths of length 50.

Method	Average Joint Entropy
Random	379.77
Straight	182.32
Ours	200.52
RIG-Tree	72.81

The random planner has the highest joint entropy value, this is likely due to the lack of path constraints and that it is sampling from the complete image, instead of a subregion. Our method performs the best out of the three path-constrained methods. We note that the paths produced by our planner and the straight-baseline planner are sample dense, that is the number of samples is close to the cost budget. The path produced by the RIG-tree is not sample dense, so there are fewer samples used in the joint entropy calculations which may be contributing to its low value.

To evaluate the planners' performance in terms of reconstruction error, we plan paths of a variety of lengths starting at eight different starting points dispersed throughout the region. Then, for nine different training images that were not used to build the model, we sample the points chosen by the planner and reconstruct the full image. We compute the reconstruction error for a given cost budget and planner by averaging the reconstruction errors for each start point and test image combination. Fig. 3 shows the average reconstruction error associated with three different planners for budgets that are multiples of between 10 and 100. The lower bound shown in the graph is the average reconstruction error per pixel when all of the data points in the test images are used to reconstruct.

We note that there is no formal relationship between joint entropy and reconstruction error, which explains the non-monotonicity of our planner's performance. The random baseline planner reaches the same average reconstruction error at budget 70 as sampling every pixel in the test images would produce. We observe that even though our planner




Fig. 3: The relationship between path length and the average reconstruction error per known pixel.

does not reach this minimum reconstruction error, it does come fairly close by budget 100.

While it is not feasible to compute the reconstruction error for the baseline planner beyond path length 100, we found that the reconstruction error for the straight path baseline did not reach the minimum reconstruction error until the budget reached 10,000.

## VI. CONCLUSIONS AND FUTURE WORK

In this paper we have considered the problem of sampling regions of a space to learn about the global state of the space by exploiting its structure. To this end, we have introduced a planner for planning cost-constrained paths through a space such that the entropy of the sampled points is bounded suboptimal. In addition we have proposed the use of generalized low-rank models for predicting the global state of a space after observations along a path have been collected.

A main limitation of our planner are memory costs associated with computing the joint entropy using a covariance matrix. We plan to explore online planning algorithms that can more efficiently exploit structure that is discovered in the space during exploration.

Currently, our evaluation is limited to a single dataset, which has a high amount of structure. When we assume Euclidean-distance path lengths and uniformly-dispersed sampling points, it has few limitations. Future work will include evaluation on some other, more challenging datasets.

## VII. ACKNOWLEDGMENT

The authors would like to thank Madeleine Udell and Robert Kleinberg for their helpful feedback.

## REFERENCES

- [1] Chenghui Cai and Silvia Ferrari. Information-driven sensor path planning by approximate cell decomposition. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(3):672–689, 2009.
- [2] Benjamin Charrow, Gregory Kahn, Sachin Patil, Sikang Liu, Ken Goldberg, Pieter Abbeel, Nathan Michael, and Vijay Kumar. Information-theoretic planning with trajectory optimization for dense 3d mapping. In *Robotics: Science and Systems*, volume 11. Rome, 2015.
- [3] Bruce Golden, Lawrence Bodin, T Doyle, and W Stewart Jr. Approximate traveling salesman algorithms. *Operations research*, 28(3-part-ii):694–711, 1980.
- [4] Geoffrey A Hollinger and Gaurav S Sukhatme. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287, 2014.
- [5] Stefan Jorgensen, Robert H Chen, Mark B Milam, and Marco Pavone. The team surviving orienteers problem: routing teams of robots in uncertain environments with survival constraints. *Autonomous Robots*, 42(4):927–952, 2018.
- [6] Chun-Wa Ko, Jon Lee, and Maurice Queyranne. An exact algorithm for maximum entropy sampling. *Operations Research*, 43(4):684–691, 1995.
- [7] Jun-young Kwak and Paul Scerri. Path planning for autonomous information collecting vehicles. In *2008 11th International Conference on Information Fusion*, pages 1–8. IEEE, 2008.
- [8] John S Langford and Kerry A Emanuel. An unmanned aircraft for dropwindsonde deployment and hurricane reconnaissance. *Bulletin of the American Meteorological Society*, 74(3):367–376, 1993.
- [9] Cindy Leung, Shoudong Huang, Ngai Kwok, and Gamini Dissanayake. Planning under uncertainty using model predictive control for information gathering. *Robotics and Autonomous Systems*, 54(11):898–910, 2006.
- [10] Kian Hsiang Low, John M Dolan, and Pradeep Khosla. Information-theoretic approach to efficient adaptive path planning for mobile robotic environmental sensing. In *Nineteenth International Conference on Automated Planning and Scheduling*, 2009.
- [11] Kian Hsiang Low, John M Dolan, and Pradeep Khosla. Active markov information-theoretic path planning for robotic environmental sensing. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 753–760. International Foundation for Autonomous Agents and Multiagent Systems, 2011.
- [12] Wenjie Lu, Guoxian Zhang, and Silvia Ferrari. An information potential approach to integrated sensor path planning and control. *IEEE Transactions on Robotics*, 30(4):919–934, 2014.
- [13] Kai-Chieh Ma, Lantao Liu, and Gaurav S Sukhatme. An information-driven and disturbance-aware planning method for long-term ocean monitoring. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2102–2108. IEEE, 2016.
- [14] Kai-Chieh Ma, Lantao Liu, Hordur K Heidarsson, and Gaurav S Sukhatme. Data-driven learning and planning for environmental sampling. *Journal of Field Robotics*, 35(5):643–661, 2018.
- [15] Joseph L Nguyen, Nicholas RJ Lawrance, Robert Fitch, and Salah Sukkarieh. Real-time path planning for long-term information gathering with an aerial glider. *Autonomous Robots*, 40(6):1017–1039, 2016.
- [16] Amarjeet Singh, Andreas Krause, and William J Kaiser. Nonmyopic adaptive informative path planning for multiple robots. In *Twenty-First International Joint Conference on Artificial Intelligence*, 2009.
- [17] Ryan N Smith, Mac Schwager, Stephen L Smith, Burton H Jones, Daniela Rus, and Gaurav S Sukhatme. Persistent ocean monitoring with underwater gliders: Adapting sampling resolution. *Journal of Field Robotics*, 28(5):714–741, 2011.
- [18] Stephen L Smith, Mac Schwager, and Daniela Rus. Persistent monitoring of changing environments using a robot with limited range sensing. In *2011 IEEE International Conference on Robotics and Automation*, pages 5448–5455. IEEE, 2011.
- [19] Madeleine Udell, Corinne Horn, Reza Zadeh, Stephen Boyd, et al. Generalized low rank models. *Foundations and Trends® in Machine Learning*, 9(1):1–118, 2016.
- [20] Junming Wang, Theodore W Sammis, Vincent P Gutschick, Mekonnen Gebremichael, Sam O Dennis, and Robert E Harrison. Review of satellite remote sensing use in forest health studies. *The Open Geography Journal*, 3(1), 2010.
- [21] Jingjin Yu, Mac Schwager, and Daniela Rus. Correlated orienteering problem and its application to informative path planning for persistent monitoring tasks. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 342–349. IEEE, 2014.
- [22] Guoxian Zhang, Silvia Ferrari, and M Qian. An information roadmap method for robotic sensor path planning. *Journal of Intelligent and Robotic Systems*, 56(1-2):69–98, 2009.
- [23] Haifeng Zhang and Yevgeniy Vorobeychik. Submodular optimization with routing constraints. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.