

# Efficient Terrain Driven Coral Coverage Using Gaussian Processes for Mosaic Synthesis

Sandeep Manjanna, Nikhil Kakodkar, Malika Meghjani, and Gregory Dudek  
School of Computer Science,  
McGill University  
email:{msandeep, nikhil, malika, dudek}@cim.mcgill.ca

**Abstract**—In this paper we address the development of efficient methods for visual mapping of outdoor environments using exploration and reward identification followed by selective visual coverage. In particular, we consider the problem of visual mapping of a shallow water coral reef to provide an environmental assay. Our approach has two stages based on two classes of sensor: bathymetric mapping and visual mapping. We use a robotic boat to collect bathymetric data using a sonar sensor for the first stage and video data using a visual sensor for the second stage. Since underwater environments have varying visibility, we use the sonar map to select regions of potential value, and we efficiently construct the bathymetric map from sparse data using a Gaussian Process model. In the second stage, we collect visual data only where there is good potential payoff, and we use a reward-driven finite-horizon model akin to a Markov Decision Process to extract the maximum amount of valuable data in the least amount of time. We show that a very small number of sonar readings suffice on a typical fringing reef. We validate and demonstrate our surveying technique using real robot in the presence of real world conditions such as wind and current. We also show that our proposed approach is suitable for visual surveying by presenting a visual collages of the reef.

**Keywords**-Gaussian Processes; Selective Coverage; Coral Mosaic;

## I. INTRODUCTION

In this paper we present an efficient technique to build a depth map of the sea-bed with sparse measurements from a depth sensor and selectively cover the hot-spots, defined by a reward function depending on the depth, to obtain a detailed visual survey of the coral reef. Taking a general perspective, our work looks at using a reward function over the plane to develop an approach to sequential data collection, potentially on a recurring basis and in the presence of disturbances (such as wind). We further assume that a (different) sensor can be used to measure this reward function and this process of reward estimation also needs to be conducted efficiently. This emphasis on efficient data collection is motivated by our interest in field robotics, and especially marine robotics, where large areas need to be covered and the efficient coverage in the face of limited power, physical demands on human supervisors, and external influences is significant.

Environment monitoring is activity of enormous and growing importance. In the marine environment, this mo-



Figure 1: The Autonomous Surface Vehicle (ASV) used in field experiments

tivated by a desire to observe factors such as the impact of sea level change or increase in temperature, acidification, and changes in fauna (such as depletion of coral). Modeling these and related phenomena calls for enormous amounts of data sometimes sampled over a long durations. Collecting such data involves risk to human and various resource constraints, and leads to challenges in terms of consistency and repeatability.

As an application context and experimental validation we consider visual mapping of shallow water coral reefs. Shallow water marine environments including wetlands and estuaries have exceptional environmental, scientific and economic value and coral reefs are the most important of all. Despite their immense importance, coral reefs in particular are dying worldwide and a first step in estimating the impact of any remediation effort is to evaluate the reef itself. Conducting reef surveys is, however, both expensive and potentially dangerous. While underwater observations remain indispensable, valuable data that is economical, repeatable, scientifically objective and effective can often be collected from the surface. Doing so implies taking repeated measurements under varying weather conditions and in a consistent manner. Using an autonomous vehicle for collecting such visual data from coral reefs is efficient and relatively risk-free, although the limitations on available power make efficient coverage particularly important (notably using an electric vehicle, instead of a gas-powered one, facilitates access to protected marine environments).

We seek to efficiently build a bathymetric model (depth map of the sea floor) using data collected from our robot with a relatively short trajectory and a limited number

of samples. We use a Gaussian process (GP) to infer a complete terrain map from these samples. We then seek to efficiently cover the most visually valuable regions of this map using a data-driven coverage algorithm. In this case, we exploit the idea that the shallowest water allows for the best measurements of the underlying reef, and is thus the most valuable. Other measurables can, in principle, be used for the reward function including visual diversity and environmental conditions. The sea floor is, of course, a complex surface especially in the neighborhood of fringing reefs which have complex 3D structures. To build these models we exploit the power of Gaussian processes regression (also known as Kriging [1]) to represent these data-driven functions, interpolate between measurements, and also provide a covariance estimate at each point that can be used to drive further data acquisition.

One of the essential features of a good coverage algorithm is to visit the hot-spots in the increasing order of their value. This becomes significant when the task is to cover a given region in limited time. This behavior of selecting the salient regions to examine first is the key feature of our selective coverage algorithm. We present a value-iteration based algorithm which covers the entire region of interest, but in a prioritized fashion. We assume an underlying distribution for the phenomenon that needs to be modeled and build an off-line trajectory to cover the high probability regions along with a concern to reduce the travel time and energy consumption. Some of the potential applications would be - collecting samples from the ocean based on the surface temperature data from static sensor nodes or temperature maps from satellites, sampling visual data of coral reefs to monitor their health and growth, and sampling atmospheric gases based on the satellite maps. Off-line planning is essential when the mission is time and precision critical. We validate and demonstrate the results on a robot boat (Fig. 1) in the presence of wind and current.

## II. RELATED WORK

Environment monitoring has been a great concern for scientists all over the world due to several factors like growing sea-levels, increase in the global temperatures, increased sea surface temperatures, and depletion of corals. Science community is very much interested to understand the phenomena of nature. But, this needs enormous amount of data sampled over a long duration of time. These environments are not very safe for humans to sample regularly. This makes it a perfect application for using autonomous agents. Collecting wide-field data, especially in a sea, is often achieved using sessile sensor nodes [2] [3]. These have an advantage in terms of consistency, but need to be replaced, often provide limited measurement density, and are rarely suitable for visual data.

Another approach to collect a high resolution data for such applications is to use an active agent to cover the region

of interest. Many methods for coverage consider uniformly interesting regions and aim at achieving a complete coverage of the region [4]. An application like search and rescue promote a complete coverage [5]. If the problem is to explore and map an unknown environment, it requires sampling the entire terrain [6] [7]. With limited resources (e.g., battery life, and time of operation), a complete coverage is impractical if the environment to be covered is large with very few hot-spots. Hollinger *et al.* [8] consider a similar problem of visiting sensor nodes deployed in the ocean using an autonomous agent as a variant of the Traveling Salesperson Problem.

Recently there is a growing interest in non-uniform coverage. Seyed *et al.* propose a coverage strategy based on space-filling curves that explore the region non-uniformly [9]. They propose a coverage tree with Hilbert-based ordering of nodes. Another interesting approach for coverage based on the model of curiosity is proposed by Girdhar *et al.* in [10]. Their information theoretic path planning technique provides paths passing through regions with higher surprise factor such that they can be used for distinguishing various terrains. However, the paths provided by such a method are highly biased by temporary variations which is good for anomaly detection but might not be suitable for persistent monitoring. We propose a selective coverage algorithm to efficiently plan a path based on the partial knowledge about the environment. Specifically, a value-iteration based algorithm which covers the most salient regions in a very efficient manner.

We propose a realistic application for our algorithm in terrain based coral monitoring. In this application, we need an underlying terrain map as a reward distribution. Several methods for modeling and representing terrains [11] exist in literature. But most of these methods do not handle spatially correlated data and do not have a convenient form to represent the uncertainty associated with the model. Gaussian Processes are very efficient in handling these cases. GP modeling is helpful in cases where parametric representation of the underlying phenomenon is not easily determined. It also effectively incorporates the uncertainty associated with the computed model. In the robotics context, several authors including Kemppainen examined sub-sampling using Gaussian Processes and considered optimal sampling strategies in terrestrial environments [12]. GP's have also proved effective in modeling the spatial distribution terrain height or other signals such as radio signal strength [13], [14]. More broadly Gaussian process regression has also proven effective in a number of other robotics contexts such as modeling the dynamics of complex systems for controller design or learning [15]. We use GP models to build the depth map of the sea floor and follow the depth map to efficiently cover the corals in the region. We also present an evaluation of GP models to build efficient terrain models with minimal sensor measurements.

### III. GAUSSIAN PROCESSES FOR DEPTH MAP

In this section we review the use of Gaussian Process (GP) models for regression, as presented in [16]. GP formulation uses a supervised learning set up in which it considers a training set  $S = \{(x_i, y_i) \mid i = 1, \dots, n\}$  is given, where  $x_i \in \mathbb{R}^D$  denotes the input locations and  $y_i \in \mathbb{R}$  denotes the sensor output at these locations. The sensor measurements  $y \in \{y_i\}_1^n$  are modelled as noisy observations of the underlying function  $f(x)$  given by  $y = f(x) + \epsilon$ , where  $\epsilon$  is assumed to be independent identically distributed Gaussian noise with zero mean and variance  $\sigma_n^2$ .

GP models the function  $f(x)$  as a multivariate Gaussian distribution, completely defined by the mean function  $m(x)$  and the covariance function  $k(x, x')$ , also known as the kernel, i.e.  $f(x) \sim \mathcal{N}(m(x), k(x, x'))$ . The training set is used to learn this model by tuning the mean and the covariance functions. Using the learned model, predictions on new inputs  $X_* = \{x_{*1}, \dots, x_{*m}\}$  is described by  $\mathcal{N}(\mu_*, \Sigma_*^2)$ , where

$$\mu_* = K(X_*, X)[K(X, X) + \sigma_n^2 I]^{-1} y \quad (1)$$

$$\Sigma_* = K(X_*, X_*) - K(X_*, X)[K(X, X) + \sigma_n^2 I]^{-1} K(X, X_*) \quad (2)$$

Here  $K$  represents the covariance matrix calculated between the corresponding inputs as defined by the covariance function. Intuitively covariance function describes the relationship between two outputs as a function over their inputs (location in our case).

### IV. EFFICIENCY OF GPs WITH COARSE DATA

Given a region of interest, we seek to survey the corals occurring within the area using the Autonomous Surface Vehicle (ASV). First stage in our approach involves conducting a quick bathymetric survey and generate a depth map by employing the GP regression. Since we are limited by factors such as robot battery our goal is to to conduct this survey with relatively shorter trajectory and fewer sample points than it would take to do a complete dense coverage. To measure the reliability of a depth map generated from fewer number of samples, we first conducted a dense survey of the region shown in Fig. 2. This lawnmower was completed in two charge cycles of the vehicle, which can be seen in Fig. 2 where the robot was brought back to the shore mid-mission to replace the batteries and the mission was resumed from where it had left off.

From this dense coverage, we collected a total of 2962 sample points using a sonar sensor running at 1Hz. Fig. 3 illustrates the performance of the GP regression technique computed as root mean squared error versus the number of training samples. For this evaluation we randomly divide the 2962 sample points into eight sets. Hence at a time there are 2562 training points and 400 test points. We then perform 8-fold cross-validation using these sets. During each fold of the



Figure 2: Dense data collected by the robot.

cross-validation we carry out the GP formulation multiple times, each time randomly choosing  $n$  training points from the training set,  $n$  starting from 100 with increments of 300 points. For each  $n$  we repeat the procedure 50 times, each time re-selecting  $n$  samples from the training set. Finally average the mean and standard deviation for every trial with  $n$  sample points over all the 8-folds. The results indicate that a sparse sampling in the range of (1200 : 2000) samples would still model the terrain well enough for our application domain.




Figure 3: MSE vs No. of samples.

### V. SELECTIVE COVERAGE

In the previous sections we addressed the problem on mapping sea floor with minimum number of sensor readings. We use depth map to achieve the coverage of the coral-heads in the selected region. We make an assumption that in a given area, the coral-heads are at shallower depths than the plain sand. We make this assumption based on our observations during field trials and inputs from the experts in marine robotics. This assumption is true when we are away from the shore line as the sea bed tends to be smooth with corals rising above the sand surface. Based on this assumption, covering the shallower regions would provide a good coverage of the corals in the selected region.

In the case of coral monitoring, we use the actual depth map as an input distribution to cover the shallower regions with a precedence, thus making sure the corals are fully

covered before the robot runs out of battery. One of the salient features of this algorithm is the use of predictions for environment dynamics, such as wind direction, wind speed, and wave currents, to generate an efficient off-line trajectory. Thus, the proposed approach adapts the trajectories according to the wind prediction inputs. Once we have the depth map of a certain region which we are interested in monitoring, the selective coverage algorithm can be used to generate coverage trajectories for regular monitoring of the coral reef.

#### A. Algorithm

In our approach, the region of interest is discretized into grid cells and each grid-cell is assigned a utility value equivalent to the integral of underlying probability distribution over that cell. In the coral coverage task, the utility value of a grid-cell is summed over the depth distribution provided by the Gaussian Processes. We assume a finite set of states  $S$  represented by these grid-cells and a finite set of actions  $A$  the agent can take at each state (grid-cell). We allow four actions in every state (grid-cell) - East, North, West, and South. We use GPS to localize the agent and achieve a guaranteed state transition with the help of a precisely tuned controller. Hence, we assume certainty in the transition between two states  $s, s' \in S$  when an action  $a \in A$  is taken, i.e., the probability of state transition  $P(s'|s, a) = 1$ . As mentioned earlier, the agent collects a reward  $R(s, a)$  for visiting a state  $s \in S$  and taking an action  $a$  in state  $s$ .

We use the *Value Iteration* algorithm to compute the best action to be taken at a given state. Value iteration is a method of computing an optimal MDP policy. It computes the optimal value of a state  $V^*(s)$ , i.e., the expected discounted sum of rewards that the agent will achieve if it starts at that state and executes the optimal policy  $\pi^*(s)$ .  $\forall s \in S$ , the optimal value function  $V^*(s)$  is defined by the following Bellman equation,

$$V^*(s) = \max_a \left( R(s, a) + \gamma \sum_{s' \in S} P(s'|s, a) V^*(s') \right), \quad (3)$$

where  $\gamma$  is a discount factor. Thus according to Eq. 3, the value of a state  $s$  is the sum of instantaneous reward and the expected discounted value of the next state, when the best available action is used. Optimal policy defines an action for every state that achieves the optimal value. Given the optimal value function for all states, optimal policy is defined by,

$$\pi^*(s) = \arg \max_a \left( R(s, a) + \gamma \sum_{s' \in S} P(s'|s, a) V^*(s') \right) \quad (4)$$

In our approach, the rewards are defined by the summation of depth predictions under that state. We clear the underlying

rewards as and when the corresponding state is visited. Thus the reward function is changing over time as the agent clears the rewards. But this is against the Markovian assumption to keep a track of visited states. Hence, in our formulation we use an one step MDP approach, where we model every state transition of the agent as MDP in a new world and compute the value function over the updated rewards of the world. Thus the convergence of the value iteration technique still holds good for every state transition. The overview of our approach is presented in the Algorithm 1.

---

#### Algorithm 1 Selective Coverage Algorithm

---

**Input:** Set of states  $S$

Set of actions  $A$

State transition probability  $P(s'|s, a)$   
 $\forall (s, s') \in S$  and  $\forall a \in A$

Reward  $R(s, a)$  for each state  $s \in S$

Discount factor  $\gamma$

Starting state  $s_1$

Convergence threshold  $R_{limit}$

Convergence threshold  $\epsilon$

Iteration threshold  $I_{limit}$

**Output:** Path  $\vec{W} = (s_1, s_2, \dots, s_n)$ , a sequence of states.

- 1:  $\forall s \in S$ ,
  - 2:     **Initialize**  $V^*(s)$ ,  $\pi^*(s)$ , and current state  $s_{cur} = s_1$
  - 3: **Repeat**
  - 4:      $\vec{W} = \mathbf{Append}(\vec{W}, s_{cur})$
  - 5:      $\forall s \in S$ ,
  - 6:          $(V^*(s), \pi^*(s)) = \mathbf{ValueIteration}(S, A, P, R, \gamma, \epsilon)$
  - 7:     **Current Action**,  $a_{cur} = \pi^*(s_{cur})$
  - 8:     **ApplyAction**  $a_{cur}$  on  $s_{cur}$  to obtain  $s_{next}$
  - 9:      $R(s_{cur}) = 0$ , Clearing the reward at  $s_{cur}$
  - 10:      $s_{cur} = s_{next}$
  - 11: **until**  $(\sum_{s \in S} R(s) < R_{limit})$  or the region is fully covered.
  - 12: **Return**  $\vec{W}$
- 

#### B. Trajectory adaptation

The proposed algorithm has the capability to adapt its output trajectory in accordance with predictions about the changes in the operating environments. In particular for coral monitoring, we use the predictions for wind speed and wind direction to generate efficient coverage trajectories. The proposed approach is very useful to generate trajectories for re-sampling the given region of interest over a period of time. We incorporate the effect of wind on the agent's trajectory by subtracting a cost term into the reward function  $R(s, a)$ . Based on our observations in the field we make an intuitive assumption that the horizontal component of wind ( $h_{wind}$ ) affects the vertical trajectories of the agent and the vertical component of wind ( $v_{wind}$ ) affects the horizontal trajectories. Hence, we subtract the weighted  $h_{wind}$  from




Figure 4: Trajectories generated by Selective Coverage algorithm over different complex reward distributions. The colorbars indicate the rewards.

rewards  $R(s, a)$  when the action  $a = \{North, South\}$ . Similarly, the weighted  $v_{wind}$  is deducted from rewards  $R(s, a)$  when the action  $a = \{East, West\}$ . The updated rewards are indicated in Eq. 5 and Eq. 6.

$$\forall a = \{North, South\} \text{ and } \forall s \in S,$$

$$R(s, a) = R(s, a) - (\epsilon * w_{speed} * \sin(w_{\theta})), \quad (5)$$

$$\forall a = \{East, West\} \text{ and } \forall s \in S,$$

$$R(s, a) = R(s, a) - (\epsilon * w_{speed} * \cos(w_{\theta})), \quad (6)$$

where  $\epsilon$  is the convergence factor for the value iteration,  $w_{speed}$  is the speed of wind in  $ms^{-1}$ , and  $w_{\theta}$  is the wind direction in radians.

### C. Simulation Results

In this section, we illustrate the output of our selective coverage technique on different underlying probability distributions. We illustrate the capability of the algorithm to generate efficient trajectories with a diverse set of underlying distributions. Fig. 4(a) shows the trajectories generated on three different kinds of distributions. As it can be seen, the trajectory covers the highest rewarding states followed by lower rewarding states. But, the trajectories generated are not greedy, i.e., as it is seen in Fig. 4(b) and Fig. 4(c), the algorithm efficiently trades-off between the rewards and the distance of the states. A near-by state with slightly lesser reward is preferred over the farther state with slightly higher rewards.

## VI. FIELD EXPERIMENTS AND RESULTS

Our motivation in this paper is to provide efficient trajectories for sampling and re-sampling the regions of interest based on depth and we examined the performance of the method in the open sea. To perform this task autonomously we provide the autonomous boat with a time and power efficient trajectory based on the depth map collect in a first

sampling path, under the supervision on human monitors on shore. We assess the performance of the coverage path generated by our algorithm with respect to three metrics - the path length, the time taken to execute the trajectory on real robot, and the power consumption by the robot motor.

### A. Setup

We conducted field experiments on the North Bellairs reef in the Caribbean Sea off the shore of Barbados. An Autonomous Surface Vehicle (ASV), shown in Fig. 1, was used to collect the depth measurement data and execute the generated coverage trajectories. The maximum operating speed of the ASV is  $1.3ms^{-1}$  and the battery life is 1.5 hours. The ASV is equipped with navigational sensor suit consisting of a global positioning unit (GPS) and an inertial measurement unit (IMU). The GPS operates at a frequency of  $5Hz$ . The ASV is also equipped with a sensor payload consisting of a single beam sonar and a water-sealed camera unit. The sonar has a maximum depth range of  $100m$  and operates at  $1Hz$ . The camera unit records videos at  $30fps$ .

The robot was operated from the shore via a wifi link. It can also be teleoperated using a RC gamepad. All the communication was relayed via a custom GUI shown in fig 5. We use the GUI to design and relay custom trajectory patterns and view important mission parameters. Figure 5 illustrates how we can use the GUI to select user defined waypoints and display the ASV trajectory, in real time, as it makes its progress. The GUI also displays the 'Home' co-ordinate and the 'Current Target' co-ordinate.

We evaluated the proposed coverage approach over a region of size  $110m \times 75m$  indicated by the rectangular box in Fig. 5.

### B. Depth Map and Coverage Trajectories

We used the sparse sonar data from the robot boat to generate a depth map of the sea floor using the GP models as described in Section III. A sparser scanning trajectory

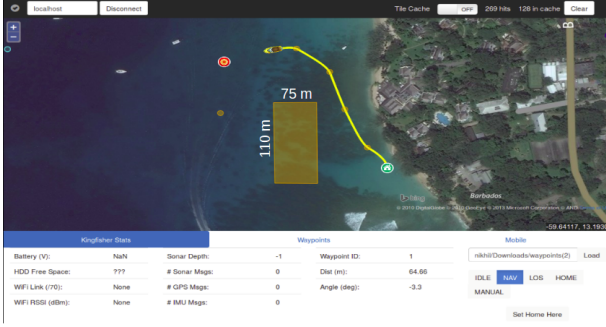


Figure 5: Graphical User Interface: Green marker indicates the 'Home' location while the red marker shows the 'Current Target Waypoint'. Each of the yellow circles shows the waypoints selected for the ASV to traverse to. The GUI also draws the trail of the robot shown in yellow.



Figure 6: Sparse lawnmower trajectory performed by the boat to collect the bathymetric data to create the depth map.

(Fig. 6) with 1340 sample points was used to generate the depth map with an acceptable error of approximately  $0.2m$  according to our analysis in Section IV. The depth map is used as an input to the selective coverage algorithm and the output trajectory overlaid on the depth map is illustrated in the Fig. 7. The colorbar in the Fig. 7 represents the depth measurements with deep red being the shallowest region and deep blue being the deepest. The coverage region displays thr2.8206ee shallower islands indicating high-rewarding regions for the selective coverage algorithm. The trajectory generated by our algorithm covers the high-rewarding islands with a priority ensuring the visible shallower regions are fully covered.

We used the generated off-line trajectory to guide the autonomous survey of the corals by the robot boat. Fig. 8 presents the actual GPS track of the robot boat performing the input trajectory from Fig. 7. The actual path followed by the boat has small discrepancies compared to the input trajectory. This is due to the non-holonomic nature of the boat and the controller parameters. Fig. 8 shows the complete GPS track of the robot in field experiment including the deployment phase on the shore.




Figure 7: Selective Coverage trajectory (in black) overlaid on the depth map. Colorbar represents the depth.



Figure 8: GPS track of the robot executing the trajectory from selective coverage. The figure also shows the deployment phase on the shore.

### C. Mosaic

We collected the visual data from the surface of the ocean with the robot boat performing the coverage pattern. We used the visual data from a transect of the boat trajectory to generate the mosaic presented in Fig. 9. Fig. 10 presents a collage of salient geo-referenced visual frames extracted from the data collected to provide a quick visualization. The trajectory is overlaid on the collage to indicate the mapping between the covered region and the geo-referenced visual frames.

### D. Performance Analysis

We compare the performance metrics (path length, trajectory execution time, and power consumption by the boat motor) of the selective coverage trajectory (Fig. 8) with the measurements from the dense lawnmower coverage (Fig. 2) to get a sense relating to task achievability with limited battery on the robot boat.

The above Table I provides a side-by-side comparison of the performance of the two coverage patterns. it illustrates

Coverage Trajectory	Path length (m)	Time (s)	Power Consumption (Wh)
Selective Coverage	1612.78	2423	33.8472
Lawnmower	3119.4	4363	67.9356

Table I: Performance measurements for trajectories from selective coverage and lawnmower.

the superior performance of the value based mechanism described in this paper and allows sampling to occur over wider regions than otherwise possible with a limited power budget. The importance of this budget is exemplified by the fact that, as mentioned in Section IV, the robot boat ran out of battery before completing the coverage of the selected region using the dense lawnmower pattern (Fig. 2). Hence, two iterations (i.e. with a fresh battery) were needed to complete a single run. The above comparison of the trajectory provided by our algorithm to the worst case lawnmower coverage shows that we can achieve the task of coral monitoring efficiently under the real world constraints, such as battery time and execution time.

## VII. CONCLUSION

In this paper we described an approach to surface-based reef monitoring using a 2-phase process of bathymetric mapping followed by visual data collection. In practice we expect the latter of these stages to be conducted repeatedly. The first phase employs a Gaussian process and we show that this allows an acceptable depth map to be computed using fewer samples than would be needed for more naive methods. The second phase we exploit this bathymetric data along with a value based sampling methodology and show this allows a complete map to be computed much more efficiently than traditional methods such as repeated linear transects. We validated this overall approach in the open sea and produced a visual coverage map that illustrated the effectiveness of the method.

In future work we plan to explore the impact of our wind model and how it related to the efficiency of the coverage pattern. This appears to be a useful aspect of our algorithm and is motivated by a real need, but a rigorous validation is both time consuming, painful to execute and outside the scope of this paper.

## REFERENCES

- [1] G. Matheron, "Principles of geostatistics," *Economic geology*, vol. 58, no. 8, pp. 1246–1266, 1963.
- [2] R. N. Smith, Y. Chao, P. P. Li, D. A. Caron, B. H. Jones, and G. S. Sukhatme, "Planning and implementing trajectories for autonomous underwater vehicles to track evolving ocean processes based on predictions from a regional ocean model," *The International Journal of Robotics Research*, p. 0278364910377243, 2010.
- [3] S. L. Nooner and W. W. Chadwick, "Volcanic inflation measured in the caldera of axial seamount: Implications for magma supply and future eruptions," *Geochemistry, Geophysics, Geosystems*, vol. 10, no. 2, 2009.
- [4] A. Xu, C. Viriyasuthee, and I. Rekleitis, "Efficient complete coverage of a known arbitrary environment with applications to aerial operations," *Autonomous Robots*, vol. 36, no. 4, pp. 365–381, 2014.
- [5] M. Meghiani, F. Shkurti, J. C. G. Higuera, A. Kalmbach, D. Whitney, and G. Dudek, "Asymmetric rendezvous search at sea," in *Computer and Robot Vision (CRV), 2014 Canadian Conference on*. IEEE, 2014, pp. 175–180.
- [6] W. Burgard, M. Moors, C. Stachniss, and F. E. Schneider, "Coordinated multi-robot exploration," *Robotics, IEEE Transactions on*, vol. 21, no. 3, pp. 376–386, 2005.
- [7] H. Choset, "Coverage for robotics—a survey of recent results," *Annals of mathematics and artificial intelligence*, vol. 31, no. 1–4, pp. 113–126, 2001.
- [8] G. A. Hollinger, S. Choudhary, P. Qarabaqi, C. Murphy, U. Mitra, G. S. Sukhatme, M. Stojanovic, H. Singh, and F. Hover, "Underwater data collection using robotic sensor networks," *Selected Areas in Communications, IEEE Journal on*, vol. 30, no. 5, pp. 899–911, 2012.
- [9] S. A. Sadat, J. Wawerla, and R. Vaughan, "Fractal trajectories for online non-uniform aerial coverage," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 2971–2976.
- [10] Y. Girdhar, D. Whitney, and G. Dudek, "Curiosity based exploration for learning terrain models," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 578–584.
- [11] M. Hugentobler, "Terrain modelling with triangle based free-form surfaces."
- [12] A. KEMPPAINEN, T. MÄKELÄ, J. HAVERINEN, and J. RÖNING, "An experimental environment for optimal spatial sampling in a multi-robot system."
- [13] B. Ferris, D. Hähnel, and D. Fox, "Gaussian processes for signal strength-based location estimation," in *In Proc. of Robotics Science and Systems*, 2006.
- [14] S. Vasudevan, F. Ramos, E. Nettleton, and H. Durrant-Whyte, "Gaussian process modeling of large-scale terrain," *Journal of Field Robotics*, vol. 26, no. 10, pp. 812–840, 2009. [Online]. Available: <http://dx.doi.org/10.1002/rob.20309>
- [15] D. Nguyen-Tuong and J. Peters, "Local gaussian process regression for real-time model-based robot control," in *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, Sept 2008, pp. 380–385.
- [16] C. E. Rasmussen, "Gaussian processes for machine learning," 2006.



Figure 9: Example mosaic from a transect of the robot trajectory.




Figure 10: Selected images from the coverage. The dotted line on the image shows the coverage trajectory of the robot. Triangle and the circle indicate the start and end points of the path respectively.