Data-driven selective sampling for marine vehicles using multi-scale paths

Sandeep Manjanna[1] and Gregory Dudek[1]

***Abstract*— This paper addresses adaptive coverage of a spatial field without prior knowledge. Our application in this paper is to cover a region of the sea surface using a robotic boat, although the algorithmic approach has wider applicability. We propose an anytime planning technique for efficient data gathering using point-sampling based on non-uniform data-driven coverage. Our goal is to sense a particular region of interest in the environment and be able to reconstruct the measured spatial field. Since there are autonomous agents involved, there is a need to consider the costs involved in terms of energy consumed and time required to finish the task. An ideal map of the scalar field requires complete coverage of the region, but can be approximated by a good sparse coverage strategy along with an efficient interpolation technique. We propose to optimize the trade off between the environmental field mapping and the costs (energy consumed, time spent, and distance traveled) associated with sensing. We present an anytime algorithm for sampling the environment adaptively by following a multi-scale path to produce a variable resolution map of the spatial field. We compare our approach to a traditional exhaustive survey approach and show that we are able to effectively represent a spatial field spending minimum energy. We present results that indicate our sampling technique gathering most informative samples with least travel. We validate our approach through simulations and test the system on real robots in the open ocean.**

I. INTRODUCTION

Our adaptive coverage planner discretely samples a scalar field over an unknown region and estimates the continuous spatial field. The goal is to generate a coverage path that maximizes the information gain in terms of interesting places to sample from, while minimizing costs (energy consumed, time spent, and distance traveled). The anytime algorithm presented in this paper is guaranteed to generate an approximating path to achieve maximum rewards with minimum energy consumption even if the algorithm is interrupted before it ends. An algorithm that can return a valid solution to a problem even if it is interrupted before it ends is referred to as an anytime algorithm [1].

One of the applications of mobile robotics is exploring and mapping environmental phenomena often represented as scalar fields (e.g. temperature). Building a persistent map of an environmental scalar field for a given region over a period of time requires continuous sensing. This problem has been addressed previously by deploying a network of static sensor nodes uniformly across the region of interest [2] or by conducting a uniform survey of the region. In general, however, environmental fields of interest are non-uniform over a given region. Persistent mapping of such fields requires the mobile sensing vehicles to estimate the variations in the field, and selectively sample measurements to maximize the rewards while minimizing costs over time.

[1]School of Computer Science, McGill University, Montreal, QC H3A0E9, Canada. {msandeep,dudek}@cim.mcgill.ca

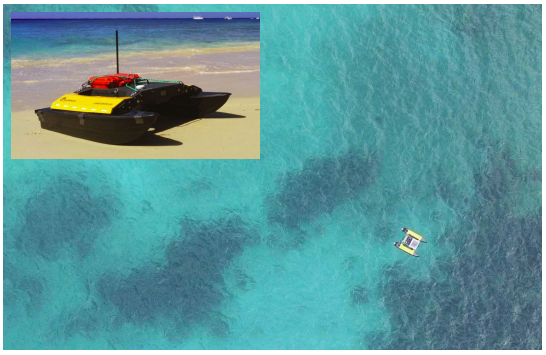

Fig. 1: Field deployment: Autonomous surface vehicle (ASV) performing Selective Sampling.

Adaptive sampling surveys of sub-regions with high local-variance has become the focus of research in robotic exploration, coverage, and mapping. The sub-regions of a spatial field which show high local-variance and associated local maxima are referred to as hotspot-regions. Spatial fields with hotspot-regions are characterized by continuous spatially correlated measurements with the hotspot-regions exhibiting extreme measurements and higher spatial variability than the rest of the field. Hence, sampling these hotspot-regions with high resolution and covering them earlier in the survey process provides a better representation of the spatial field. Our goal in this paper is to sample such environments adaptively to produce a varied-resolution maps of such spatial field and we do this in the domain of ocean surface sampling. The video [3] presents an overview of our approach.

A. Motivational Applications

Mobile sensing platforms in the form of unmanned ground vehicles (UGV), unmanned aerial vehicles (UAV), autonomous surface vehicles (ASV; Fig. 1) and autonomous underwater vehicles (AUV) can actively respond to changes in unknown environments and dynamically actuate themselves to adaptively sense the area of interest. Some of the example applications that motivate us are as listed below,

- Environmental Sensing such as monitoring of ocean phenomena (plankton blooms, upwellings, contamination) [4][5], coral reefs [6], forest ecosystems [7], rare species [8], or pollution [9].
- Terrestrial Exploration such as Antarctic meteorite search [10]; and prospecting for mineral deposits [11]; or localized methane sources on Mars [12].

- Search and Rescue missions in forests and open-water environments [13].

In this paper, we consider the task of effectively sampling the visual data of a coastal region which is known to contain coral reef outcrops and build an image mosaic to monitor the health of the reefs (ideally over successive years). The exact distribution of living coral is unknown, though we do know that the reef is not continuous over the entire region. An exhaustive survey of this area with our ASV would spend the majority of its time collecting images in which coral is not observed, wasting valuable battery power for data which is not useful for our goal of building image mosaics of the reef. In this application, the coral heads are hotspot-regions which need high resolution coverage compared to regions covered with sand. Hence, our selective coverage approach fits well for this application. We are interested in collecting visual data of shallow-water coral reefs (generally within 15m of the surface) from the ocean surface using an autonomous surface vehicle (ASV). Since we are sampling from the surface, good visual data is available at the shallower regions. Hence in this particular application sampling at shallower regions can be formalized as providing a higher reward.

B. Contributions

In prior work, we presented an algorithm that selectively samples a region with similar objectives at a single resolution [6]. In that work, we presuppose a coarse prior model of the phenomenon of interest in advance of the sampling process and thus a sampling path can be precomputed offline. This paper, in contrast, uses an online anytime algorithm and avoids the need for a prior. Our previous approach is useful when there is partial prior knowledge about the environmental field, for example from aerial surveillance. One of the limitations of our previous approach is its relatively expensive computations. In addition, on-site data sometimes proved to be at odds with observations from aerial surveys (for example if a promising region proved to contain only dead coral or other misleading visual cues), so real-time adaptation is highly desirable. The key contributions of the technique presented in this paper are that it is not being dependent on prior knowledge, that it exploits variable resolution sampling by planning multi-scale paths, and that it is relatively computationally efficient making the algorithm feasible for real-time deployment on low power devices, and that it provides anytime better result even if its interrupted in the middle. We also provide new performance data supporting these claims based on both simulation and experiments at sea.

C. Background

One is required to exhaustively sample the area as defined by the limits of the sensor to be able to completely represent a partially observable region [14], [15], [2]. The traditional approach to covering a partially observable, obstacle-free region is to employ a boustrophedon path [16]. The *boustrophedon* or *lawnmower path* is the approach a farmer takes when using an ox to plow a field, making back and forth straight passes over the region in alternating directions until the area is fully observed. We refer to the boustrophdeon approach as a metric for measuring the performance of our sampling approach.

With our approach we trade off completeness for efficiency. This notion of incomplete coverage is to be sufficient when the underlying phenomenon being sampled is band-limited, which is to say sufficiently smooth. Even for rapidly varying phenomena, sub-sampling can be effective when the sample points are correctly selected. It has been well established that for low-pass multi-band signals, uniform sampling can be inefficient and sampling rates far below the Nyquist rate can still be information preserving [17], [18]. This is the key guiding principle behind active sampling and the present work. Recently there has been a growing interest in non-uniform coverage [19], [20]. Seyed *et al.* propose a coverage strategy based on space-filling curves that explore the region non-uniformly [21]. They propose a coverage tree with Hilbert-based ordering of nodes. We have previously demonstrated an algorithm to selectively cover a region based on the underlying reward distribution [6], [13]. This technique, however, requires prior knowledge about the underlying distribution of the field which is not available in the proposed approach.

Adaptive sampling refers to strategies in which the procedure for selecting locations to be included in robot paths depends on the sampling data observed during exploration [19]. When the environmental phenomena are smoothly varying without any local-maxima peaks, non-adaptive strategies are known to perform well [22]. However, if the environment contains peaks with high local-variance, adaptive sampling can exploit the clustering phenomena to map the environmental field more accurately than non-adaptive sampling. Low et al. present an adaptive multi-robot exploration strategy that can perform both wide-area coverage and adaptive sampling of high local-variance regions using non-myopic path planning [19]. A key feature of this approach is in covering the entire adaptivity spectrum, thus allowing strategies of varying adaptivity to be formed and theoretically analyzed in their performance.

In contrast to random exploration of the environmental field [23], directed exploration selects robot paths to observe regions of high uncertainty. Directed exploration strategies that focus on feature sampling expect areas of high uncertainty to contain highly-varying measurements [11][24]. Another approach to seek the trade-off between cost and information is to use down-sampling. Compressive sensing, a recently developed down sampling and reconstructing method yielding a sub-Nyquist sampling criterion, uses condensed linear measurements for reconstruction under a sparse domain without losing useful information [25][26][27].

II. Selective Sampling

A. Overview

We assume that measurements over the region of interest have non-uniform utility and thus we collect reward in regions of maximum value, in particular points from which

we can best infer the structure of high-value areas of the scalar field of interest. In terms of implementation, we do this by imposing a uniform grid over the region of interest where each grid-cell is initially assigned a utility value equivalent to the integral of the underlying prior over that specific grid-cell. We preferentially sample grid cells that provide a good trade-off of information and accessibility. Thus our goal is to harvest most informative samples (rewards) with least expenditure of energy.

We need a path for an agent to cover hotspot-regions with high resolution and remaining region with lower resolution. We apply an efficient method of computing an optimal MDP (Markov Decision Process) policy called value iteration to compute the best action in a given state. Optimal value of a state $V^*(x)$, $\forall x \in X$, is the discounted sum of instantaneous rewards $R(s,a)$ and the expected discounted value of the next state $V^*(x')$, when the best available action ($a \in A$) is used. It is defined by the Bellman equation [28],

$$V^*(x) = \max_a \left(R(s,a) + \gamma \sum_{x' \in X} P(x'|x,a) V^*(x') \right), \quad (1)$$

where γ is the discount factor. Optimal policy ($\pi^*(s)$) defines an action for every state that achieves the optimal value. Given the optimal value function for all states, optimal policy is defined by,

$$\pi^*(s) = arg\max_a \left(R(s,a) + \gamma \sum_{s' \in S} P(s'|s,a) V^*(s') \right) \quad (2)$$

In our approach (Algorithm 1), we clear the rewards when the corresponding state is visited (as shown in Fig.2). Thus the reward function is changing over time as the agent clears the rewards. Even though this violates the Markov assumption over the entire coverage task, every state transition still holds good to be formulated as a one step MDP, where every state transition of the agent is modeled as MDP in a new world and the value function is computed over the updated rewards of the world. Thus the convergence of the value iteration technique still holds good for every state transition.

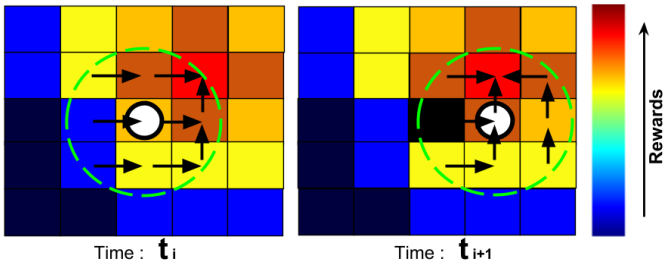

Fig. 2: The figure depicts one step of an agent (white circle) running selective sampling algorithm from time t_i to t_{i+1}. The colors in the grid-cells (states in our case) represent the rewards and arrow marks represent the optimal policy for each state. Green dotted circle represents the subset of states considered for value iteration.

This method of re-evaluating the utilities of all states iteratively for every step becomes computationally expensive and unrealistic to run in real-time. Hence, we apply an approximation for updating the utilities of the states. Value at any state is influenced more by its neighborhood than by a state that is very distant from it. Hence, for a one-step transition we achieve good utility updates by just considering a neighborhood subset (shown with green dotted circle in Fig.2) of the state space ($S_n \subset S$) instead of complete state space (S).

Algorithm 1 Online Selective Sampling Algorithm

Input: No. of transects n_T
Set of all locations X
Hyper-parameters for GP:
$hyp.Mean$ $m(x)\ \forall x \in X$
$hyp.Cov$ $k(x,x')\ \forall (x,x') \in X$
$hyp.lik$ l
Set of states S
Set of actions A
State transition probability $P(s'|s,a)$
$\forall (s,s') \in S$ and $\forall a \in A$
Discount factor γ
Starting state s_1
Sub-region size sub_R
Reward threshold R_{limit}
Convergence threshold ϵ
Output: Path $\vec{W} = (s_1, s_2, ...s_n)$, a sequence of states.
Predicted field Value $f(x)$ for each location $x \in X$
Predicted field Variance $Var(x)$ for each location $x \in X$

1: $start_T = \frac{size(S)}{n_T}$
2: Start sparse lawnmower at $start_T$ for n_T
3: **Gaussian Estimation:** $\forall x \in X$
4: $f(x) \sim \mathcal{GP}(m(x), k(x,x'))$
5: $R(s) = f(s) \forall s \in S$
6: $\forall s \in S$,
7: **Initialize** $V^*(s)$, $\pi^*(s)$, and current state $s_{cur} = s_1$
8: **Repeat**
9: S_{sub} is the **sub-region** of size sub_R around s_{cur}
10: $\vec{W} =$ **Append**$(\vec{W}, s_{cur})$
11: $\forall s \in S$ and **If** $(s \in S_{sub})$,
12: $(V^*(s), \pi^*(s))$ =**ValueIteration**$(S_{sub}, A, P, R, \gamma, \epsilon)$
13: Current Action, $a_{cur} = \pi^*(s_{cur})$
14: **ApplyAction** a_{cur} on s_{cur} to obtain s_{next}
15: $R(s_{cur}) = 0$, Clearing the reward at s_{cur}
16: $s_{cur} = s_{next}$
17: **Gaussian Estimation:** $\forall x \in X$
18: $f(x) \sim \mathcal{GP}(m(x), k(x,x'))$
19: $R(s) = f(s) \forall s \in S$
20: **until** $(\sum_{s \in S} R(s) < R_{limit})$ **or** the region is fully covered.
21: **Return** $\vec{W}$, $f(x)$, and $Var(x)$

B. Key properties

As discussed earlier, one of the key features of this selective sampling algorithm is its ability to provide high-rewarding path even when the algorithm is interrupted in the middle. Thus the algorithm is expected to find better and

better solutions the more time it keeps running. This quality of an anytime algorithm is very essential in robotic surveys considering the battery limitations and other challenges in the dynamic environments where such surveys are necessary.

Variable-resolution sampling is important when there is significant spatial variation on the underlying field being sampled, leading to an exploitable non-uniformity. In practice, this occurs when the area to be surveyed is large and the features that need to be sampled are concentrated at a few hotspot-regions. The utility of this approach is magnified when performing repeated mapping of a region over time for sampling a dynamically varying quantity for which a model can be inferred. The multi-scale paths with variable-resolution sampling of the selective sampler are illustrated in Fig. 3(c). The sample points (represented with white dots) are dense close to the high-reward locations and become sparse as the robot moves towards lower-reward regions.

III. Efficient reconstruction

Efficient coverage requires an adequate method to reconstruct the spatial field with as sparse samples as possible. Field reconstruction in an autonomous survey setup needs to deal with noisy measurements, unevenly distributed observations, and fill small gaps in the areas with high confidence while assigning higher predictive uncertainty in sparsely sampled areas. One of the techniques to naturally handle these necessities is Gaussian Process (GP) models. GP models also have the advantage of not assuming a fixed discretization of the space and of additionally providing predictive uncertainties. The explicit model of uncertainty that a GP provides has led to their successful application in a wide range of other robotic applications. Hence, we model the distribution of valuable/interesting locations in the world as a Gaussian Process: a probabilistic representation of a scalar field [29]. In this modeling of the spatial phenomena, the value of the field $f(x)$ at each location x is estimated by a Gaussian distribution with mean $m(x)$ and a covariance function or the kernel $k(x, x')$:

$$f(x) \sim \mathcal{GP}(m(x), k(x, x')) \tag{3}$$

The assumptions about the smoothness of the estimated field are incorporated in the covariance function. Input to the model is a set of noisy observations $\{(x_i^*, f_i^*)|i = 1, 2, ..., n\}$, where $f_i^* = f(x_i^*) + \epsilon_i$ and $\epsilon_i \sim \mathcal{N}(0, \sigma_n^2)$. A Gaussian measurement noise with zero mean is assumed over the observations. The predicted value for f at new location x is estimated as a Gaussian distribution:

$$f'(x) = \mathbf{k}(x)^T (K + \sigma_n^2 I)^{-1} \mathbf{f}^* \tag{4}$$

$$Var[f'] = k(x, x) - \mathbf{k}(x)^T (K + \sigma_n^2 I)^{-1} \mathbf{k}(x) \tag{5}$$

where $\mathbf{k}(x) = [k(x_1, x), .., k(x_n, x)]^T$, $K = [k(x, x')]_{x,x' \in \{x_i^*\}}$, and $\mathbf{f}^* = [f_1^*, f_2^*, .., f_n^*]^T$. The form of the kernel function $k(x, x')$ needs to be chosen such that it reflects the assumptions about the scalar field. For our experiments, we use the *radial basis kernel function* with the form,

$$k(x, x') = exp\left(-\frac{||x - x'||^2}{2l^2}\right). \tag{6}$$

The characteristic length-scale parameter l specifies how quickly the value of the computed scalar field becomes uncorrelated with distance. The global noise variance σ_n^2, and the length-scale l are known as the hyperparameters of the process and represented as $\Theta = (\sigma_n, l)$. To best represent the underlying data, the hyperparameters Θ need to be adapted.

One of the limitations of standard GP framework is the assumption of constant length-scale over the whole input space. Intuitively, the length-scales describe the area in which observations strongly influence each other [30]. For environmental fields with or without local-extrema, we use locally varying length-scales to account for different sample densities. In our case, where we sample spatial fields with varying resolutions, high-resolution at high local-variance regions and lower resolution at low local-variance regions, we use low length-scale for regions with high sampling density and higher length-scale when the sampling density is low.

IV. Simulations and Results

We evaluate our approach in simulation over synthetically generated spatial data and also on ground truth data gathered from field observations. Fig. 3(a) depicts the synthetic field generated by sampling three different distributions and normalizing the values. The synthetic spatial field simulates the distance between two points in the map and thus the robot simulation can simulate the actual distance traveled. We used a 4-transect sparse lawn mover with randomly chosen starting points to generate the initial reward distribution.

We validate our sampling approach by comparing the spatial maps generated by measurements sampled by selective sampling technique to the one generated by sampling on a traditional boustrophedon or lawnmower path. Fig. 3(b) and Fig. 3(c) are the maps generated by lawnmower path and our selective sampling approach respectively, after traveling a total of 1500 meters in simulation. The white markers indicate the sampling points. We quantify these maps, which indicate how well the underlying distribution is estimated, by finding the mean squared error (MSE) of the map as compared to the synthetic ground-truth distribution (Fig. 3(a)).

The statistical evaluation of our algorithm is done by running 30 simulations with varying start-state for the autonomous agent. We use three metrics to assess the performance of our approach: 1. Mean squared error (mse) between the interpolated map and the ground truth map, 2. Percentage of area covered or fraction of grid-cells visited, and 3. Percentage of value gathered or reward harvested. Fig. 4 indicates that the selective sampler achieves better reconstruction of the map with a minimal travel compared to the exhaustive sampler.

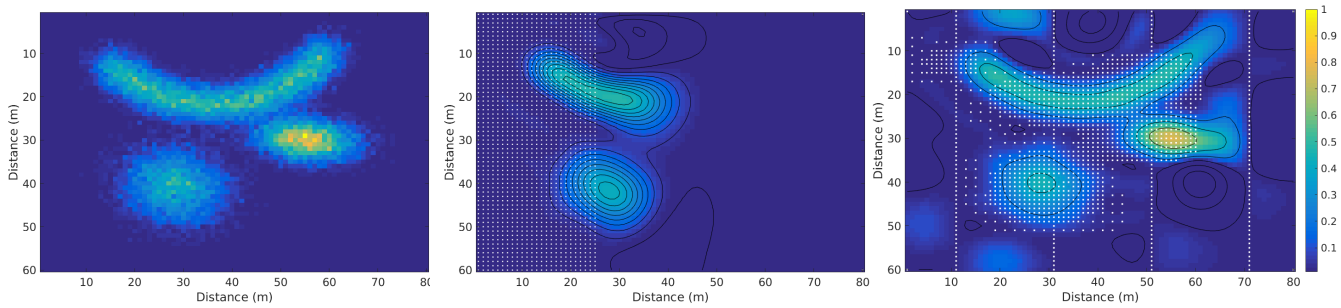

Fig. 3: (a) Synthetic ground truth spatial field - The spatial field is simulated by the distribution shown here. The values of the field are between 0 and 1 as indicated by the color bar. Both lawnmower coverage (b) and selective sampling coverage (c) are interrupted after a travel of 1500m and the spatial field map is generated using their samples. The corner samples in (c) are included to simplify the boundary implementation.

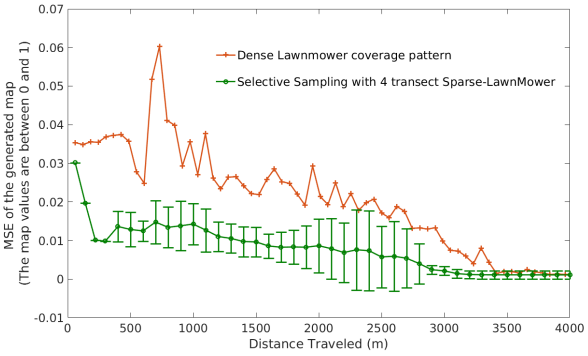

Fig. 4: Mean squared error in the generated spatial map vs. distance traveled (m). Error bars indicate two-sided standard deviation over 30 trials.

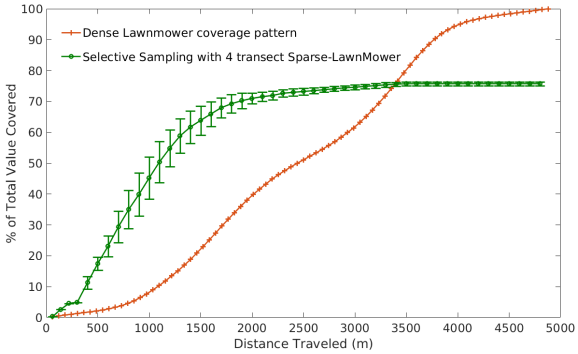

Fig. 6: Percentage (%) of rewards gathered vs. distance traveled (m). Error bars indicate two-sided standard deviation over 30 trials. Our algorithm generates paths that sample value-rich locations as quickly as possible.

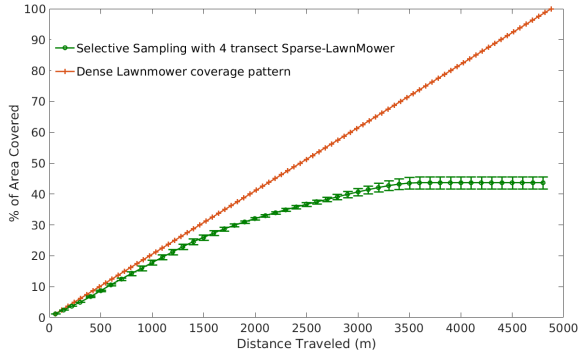

Fig. 5: Percentage (%) of area covered against distance traveled (m). Error bars indicate two-sided standard deviation over 30 trials. Our model does not visit every location in the spatial world due to varied sampling resolution.

Fig. 5 presents the plot of percentage of area covered. Here, the selective sampler does not visit 50% of the grid-cells but still collect very high rewards pretty soon in the survey as seen in Fig. 6. Lesser number of grid-cells visited imply lesser energy spent on travel. Thus our goal to harvest most informative samples (rewards) with least expenditure of energy is achieved by the selective sampler.

V. Field Experiments and Results

We validated our approach over a shallow region known to have several coral outcrops in the Folkestone Marine Reserve, Holetown, Barbados (Fig. 7(a)). We used an ASV (Fig. 1) equipped with depth pinger that measures the depth of the sea floor with acoustic pings at a frequency of $1Hz$. The ASV also has GPS ($5Hz$) for localization and a camera recording frames at a rate of $10Hz$. The surface vehicle was set to operate at half of its maximum speed (0.65 ms^{-1}).

To build a baseline, we initially sampled the region extensively (as seen in Fig. 7(b)) and built a depth map of the sea floor which is utilized as the ground truth (Fig. 7(d)) to evaluate selective sampling approach. Fig. 7(b) and (c) are the real-time screen shots of the robot's graphical interface when the robot is performing respective paths. We are interested in effectively sampling the visual data of the coral reefs and build an image mosaic to monitor the health of the reefs over years. The selective sampling algorithm generated a sampling path based on the interpolated map from sparse lawnmower samples (depicted by black trail in Fig. 7(e)). As we wanted high resolution sampling of the shallow regions, the algorithm generated sampling path that covers all the shallower regions, however most shallow

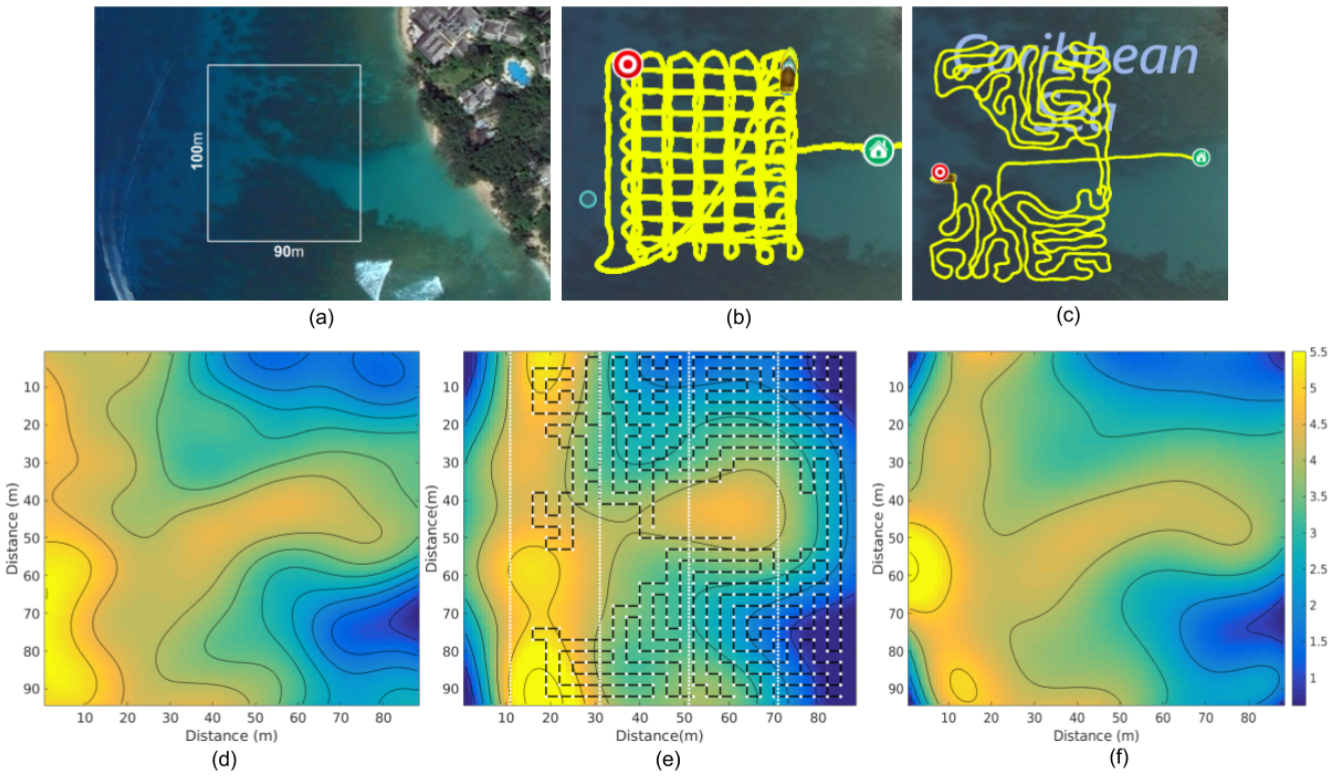

Fig. 7: (a) Experiment site in Barbados with region of interest(white box). (b) Dense lawnmower path and (c) Selective sampling path executed by the ASV (Red blob represents goal and Green represents home). (d) Depth map of sea floor generated with dense lawnmower samples. (e) Selective sampling path (black) and samples (white dots). (f) Map generated using the samples from Selective Sampling path. The color bar indicates the depth of sea floor (in meters) from the surface.

regions get higher priority. Fig. 7(c) presents the actual path followed by the ASV. The samples collected over this path of length $1800m$ were used to generate the spatial map presented in Fig. 7(f). The mean squared error between the interpolated depth map and ground truth depth map was found to be $0.1587m^2$.

We further evaluated our algorithm by running 10 trials of simulations on the real data (Fig. 7(d)) with randomly chosen starting locations for selective sampling algorithm. The Fig. 8 compares the mean squared error of the depth map generated with samples from selective sampling technique and dense lawnmower sampling path. The observations are consistent with the results from synthetic data.

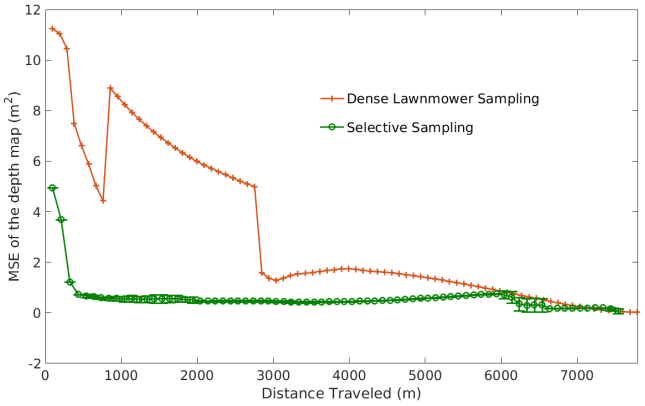

Fig. 8: Mean squared error of the depth map (m^2) vs. distance traveled (m). Error bars indicate two-sided standard deviation over 10 trials. The error bars are not visible due to very small values compared to the range of y-axis.

VI. CONCLUSIONS

In this paper we presented an anytime algorithm for active data-driven sampling of a spatial field and illustrated its deployment on an autonomous surface vehicle (i.e. a robot boat). We discovered that the algorithm was efficient and lead to productive sampling paths that led to good approximations of the phenomenon being modeled with much lower data-acquisition costs than alternative methods that have been used in practice. We presented statistically significant results, both in simulations and real robot experiments, to establish that selective sampling algorithm presented in this paper is capable of gathering most informative samples with least expenditure of energy.

A limitation of our approach is that it fails to take into account estimates of the available time horizon to complete the task, which we plan to incorporate in the future. In addition, our current deployment is a proof of concept for a more robust and high-endurance vehicle we expect to use in the future. One of the exciting directions for future work is to model the dynamics of the changing spatial field and including those dynamics into the planning of sampling paths. In the near future we plan to implement this algorithm on a multi-robot system to address the challenges in active sampling, exploration, and search and rescue.

ACKNOWLEDGMENT

We would like to thank the the Natural Sciences and Engineering Research Council (NSERC) for their support through the Canadian Field Robotics Network (NCFRN). We also appreciate our team members in the Mobile Robotics Lab at McGill University for technical discussions and assistance in field experiments.

REFERENCES

[1] S. Zilberstein, "Using anytime algorithms in intelligent systems," *AI magazine*, vol. 17, no. 3, p. 73, 1996.

[2] A. Howard, M. J. Matarić, and G. S. Sukhatme, "Mobile sensor network deployment using potential fields: A distributed, scalable solution to the area coverage problem," in *Distributed Autonomous Robotic Systems 5*. Springer, 2002, pp. 299–308.

[3] S. Manjanna and G. Dudek. Video for data driven selective sampling. Mobile Robotics Lab, McGill University. [Online]. Available: https://youtu.be/QtnzWgg9adM

[4] E. Fiorelli, N. E. Leonard, P. Bhatta, D. A. Paley, R. Bachmayer, and D. M. Fratantoni, "Multi-auv control and adaptive sampling in monterey bay," *IEEE Journal of Oceanic Engineering*, vol. 31, no. 4, pp. 935–948, 2006.

[5] N. E. Leonard, D. A. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, and R. E. Davis, "Collective motion, sensor networks, and ocean sampling," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 48–74, 2007.

[6] S. Manjanna, N. Kakodkar, M. Meghjani, and G. Dudek, "Efficient terrain driven coral coverage using gaussian processes for mosaic synthesis," in *CRV '16: Proceedings of the 2016 International Conference on Computer and Robot Vision*. IEEE Computer Society, June 2016.

[7] M. A. Batalin, M. Rahimi, Y. Yu, D. Liu, A. Kansal, G. S. Sukhatme, W. J. Kaiser, M. Hansen, G. J. Pottie, M. Srivastava, *et al.*, "Call and response: experiments in sampling the environment," in *Proceedings of the 2nd international conference on Embedded networked sensor systems*. ACM, 2004, pp. 25–38.

[8] W. Thompson, *Sampling rare or elusive species: concepts, designs, and techniques for estimating population parameters*. Island Press, 2013.

[9] H. Chang, A. Q. Fu, N. D. Le, and J. V. Zidek, "Designing environmental monitoring networks to measure extremes," *Environmental and Ecological Statistics*, vol. 14, no. 3, pp. 301–321, 2007.

[10] D. S. Apostolopoulos, M. D. Wagner, B. N. Shamah, L. Pedersen, K. Shillcutt, and W. L. Whittaker, "Technology and field demonstration of robotic search for antarctic meteorites," *The International Journal of Robotics Research*, vol. 19, no. 11, pp. 1015–1032, 2000.

[11] K. H. Low, G. J. Gordon, J. M. Dolan, and P. Khosla, "Adaptive sampling for multi-robot wide-area exploration," in *Proceedings 2007 IEEE International Conference on Robotics and Automation*. IEEE, 2007, pp. 755–760.

[12] V. Formisano, S. Atreya, T. Encrenaz, N. Ignatiev, and M. Giuranna, "Detection of methane in the atmosphere of mars," *Science*, vol. 306, no. 5702, pp. 1758–1761, 2004.

[13] M. Meghjani, S. Manjanna, and G. Dudek, "Multi-target search strategies," in *Safety, Security, and Rescue Robotics (SSRR), 2016 IEEE International Symposium on*. IEEE, 2016, pp. 328–333.

[14] W. Burgard, M. Moors, C. Stachniss, and F. E. Schneider, "Coordinated multi-robot exploration," *IEEE Transactions on Robotics*, vol. 21, no. 3, pp. 376–386, 2005.

[15] 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.

[16] H. Choset and P. Pignon, "Coverage path planning: The boustrophedon cellular decomposition," in *In International Conference on Field and Service Robotics*, 1997.

[17] R. Venkataramani and Y. Bresler, "Perfect reconstruction formulas and bounds on aliasing error in sub-nyquist nonuniform sampling of multiband signals," *IEEE Transactions on Information Theory*, vol. 46, no. 6, pp. 2173–2183, 2000.

[18] H. Landau, "Necessary density conditions for sampling and interpolation of certain entire functions," *Acta Mathematica*, vol. 117, no. 1, pp. 37–52, 1967.

[19] K. H. Low, J. M. Dolan, and P. Khosla, "Adaptive multi-robot wide-area exploration and mapping," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 23–30.

[20] M. Rahimi, M. Hansen, W. J. Kaiser, G. S. Sukhatme, and D. Estrin, "Adaptive sampling for environmental field estimation using robotic sensors," in *Intelligent Robots and Systems, 2005.(IROS 2005). 2005 IEEE/RSJ International Conference on*. IEEE, 2005, pp. 3692–3698.

[21] S. A. Sadat, J. Wawerla, and R. Vaughan, "Fractal trajectories for online non-uniform aerial coverage," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 2971–2976.

[22] A. Singh, R. Nowak, and P. Ramanathan, "Active learning for adaptive mobile sensing networks," in *Proceedings of the 5th international conference on Information processing in sensor networks*. ACM, 2006, pp. 60–68.

[23] R. McCartney and H. Sun, "Sampling and estimation by multiple robots," in *MultiAgent Systems, 2000. Proceedings. Fourth International Conference on*. IEEE, 2000, pp. 415–416.

[24] A. Singh, A. Krause, C. Guestrin, W. J. Kaiser, and M. A. Batalin, "Efficient planning of informative paths for multiple robots." in *IJCAI*, vol. 7, 2007, pp. 2204–2211.

[25] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE signal processing magazine*, vol. 25, no. 2, pp. 21–30, 2008.

[26] D. L. Donoho, M. Elad, and V. N. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Transactions on information theory*, vol. 52, no. 1, pp. 6–18, 2006.

[27] S. Huang, *Adaptive sampling with mobile sensor networks*. Michigan Technological University, 2012.

[28] R. Bellman, "Dynamic programming princeton university press," *Princeton, NJ*, 1957.

[29] C. E. Rasmussen and C. K. Williams, *Gaussian processes for machine learning*. MIT press Cambridge, 2006, vol. 1.

[30] C. Plagemann, S. Mischke, S. Prentice, K. Kersting, N. Roy, and W. Burgard, "Learning predictive terrain models for legged robot locomotion," in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2008, pp. 3545–3552.