

Collaborative Sampling Using Heterogeneous Marine Robots Driven by Visual Cues

Sandeep Manjanna[†], Johanna Hansen[†], Alberto Quattrini Li², Ioannis Rekleitis², Gregory Dudek[†]

Abstract—This paper addresses distributed data sampling in marine environments using robotic devices. We present a method to strategically sample locally observable features using two classes of sensor platforms. Our system consists of a sophisticated autonomous surface vehicle (ASV) which strategically samples based on information provided by a team of inexpensive sensor nodes. The sensor nodes effectively extend the observational capabilities of the vehicle by capturing georeferenced samples from disparate and moving points across the region. The ASV uses this information, along with its own observations, to plan a path so as to sample points which it expects to be particularly informative. We compare our approach to a traditional exhaustive survey approach and show that we are able to effectively represent a region with less energy expenditure. We validate our approach through simulations and test the system on real robots in field.

Keywords-environmental sampling; exploration; mapping; sensor nodes; autonomous surface vehicle (ASV); coral monitoring;

I. INTRODUCTION

We propose a system for adaptive sampling of an unknown marine environment with an autonomous surface vehicle (ASV) which utilizes information gathered by dispersed sensor nodes. Data-independent, exhaustive coverage of a region can be unnecessarily time consuming in cases where important features are spatially localized. Many features in marine systems which are of scientific importance are indeed localized including coral and sediment distribution, salinity, or algal concentration in confluence zones. Our approach to effective sampling revolves around the premise that we can inexpensively gather low-quality data across a broad area and synchronously transfer that knowledge to a sophisticated vehicle which is capable of directed sampling. Our strategy allows the surface vehicle to reduce its energy expenditure when compared to exhaustive coverage by selectively sampling areas which are important to its goal with only a marginal increase in hardware expenditure. This is analogous to a journalist who screens tips from the public about current events, choosing only interesting stories to cover in detail.

Coral reefs are a critical part of the marine ecosystem and support a rich diversity of life with consequent economic value and social amenity. However, sea surface temperatures have increased over the past few decades, resulting in widespread coral bleaching at an increasing rate. In this paper, we consider the task of effectively sampling the visual

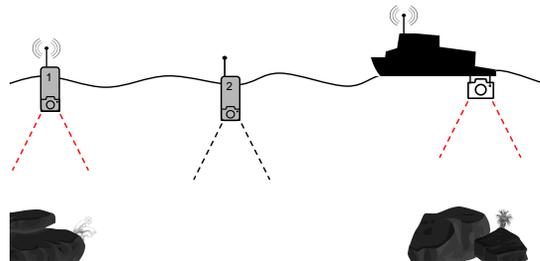


Figure 1: System overview which depicts a deployment scenario with two drifting sensor nodes and an autonomous surface vehicle (ASV). One floating sensor (drifter 1) is floating over an interesting area and reports this location to the boat via wireless communication. Another (drifter 2) is floating over an area with no interesting features does not report any location. The surface vehicle is currently sampling an interesting region, but when it finishes, it will consider the information provided by its team of drifters to choose the next optimal coverage area based on the information as described in Section III.

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 over years. The exact distribution of 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 spends the majority of 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 reefs. Having a team of heterogeneous platforms operate together allows us to reap the benefits of both the platforms - drifting sensor nodes (Figure 2(a)) which are low-cost, passively floating, less power consuming drifters, and an autonomous surface vehicle (Figure. 2(b)) which is more controllable and loaded with high quality sensors.

After a brief overview of related work in Section I-A, we describe the components of our system in Section I-B. In Section II, we discuss the extraction and scoring of image feature and then discuss how we use this information to choose areas to sample in Section III. Experimental results of simulations and field trials are presented in Sections IV and V followed by future work and conclusions.

A. Related Work

Monitoring of marine environments is important for understanding natural and human processes, but it is often

[†]School of Computer Science, McGill University ² Computer Science & Engineering Department, University of South Carolina,

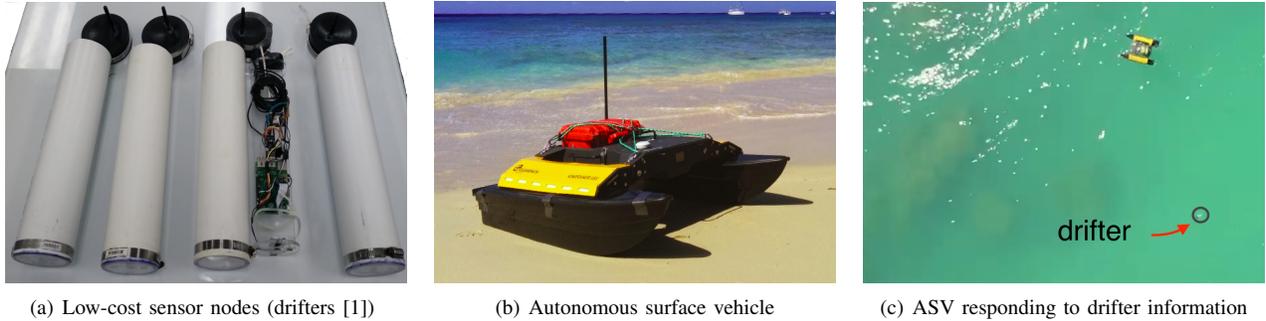


Figure 2: Hardware used in our field experiments over a region with coral outcrops.

costly and/or labor intensive to collect quality data. Many factors like fauna distribution, oxygen concentration, or coral bleaching are only locally observable and require sensors or humans to be deployed in harsh environments at or below the water surface.

As sensing platforms become more accessible, many scientists have turned to autonomous devices to collect high-resolution data in-situ. Employing a team of heterogeneous robots was shown to have distinct advantages in environmental monitoring [2]. There are generally two classes of autonomous platforms for performing these surveys. On the high-end, surveys are performed by robotic vehicles which either exhaustively survey a region [3, 4] or selectively explore a region based on local information [5]. On the low-cost end of the spectrum, numerous sensor nodes are deployed to sample points in a region, usually without actuation. Sensor nodes such as those presented in [6] are fixed to a location on the seafloor to monitor environmental changes. Others, such as the *drifter* used in our experiment [1, 7, 8], are moved about by external forces. Robotic vehicles typically have large power and computational requirements while sensor nodes tend to be designed to perform limited operations for a long period of time. Our approach seeks to combine the appealing traits of these two sampling platforms. We seek to efficiently utilize the controllability of our sophisticated robotic platform by employing information gathered by affordable, but uncontrollable sensor nodes.

To completely represent a partially observable region, one is required to exhaustively sample the area as defined by the limits of the sensor [9, 10, 11]. The traditional approach to covering a partially observable, obstacle-free region is to employ a boustrophedon path [12]. 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 boustrophedon approach as a metric for measuring the performance of our informed coverage approach.

We realize that our approach surrenders the guarantee of successfully finding all features in the region by not

surveying exhaustively, but maintain that this is a reasonable trade-off for surveys where important features are spatially localized forming hot-spots of information. In short, we trade off completeness for efficiency. Recently there is a growing interest in non-uniform coverage [13, 14]. Seyed *et al.* propose a coverage strategy based on space-filling curves that explore the region non-uniformly [15]. They propose a coverage tree with Hilbert-based ordering of nodes. We have previously demonstrated an anytime algorithm to selectively cover a region based on the underlying reward distribution [16]. This technique, however, requires prior knowledge about the underlying distribution of the field which is not available in the proposed approach. In related work it was shown that even an approximate model of a drifting sensor’s motion could enable an autonomous agent to rendezvous for data transfer [17].

A similar problem of visiting sensor nodes deployed in the ocean using an autonomous agent as a variant of the Traveling Salesman Problem has been considered by several researchers [18, 19]. Our problem is similar if we consider a node to be a point where a drifter made an observation, but we do not consider visiting all of these observation points. Instead we consider only those points which we expect to significantly increase our understanding about the region.

B. Overview of Our Approach

Our proposed method employs a sophisticated vehicle which is able to strategically sample points where it expects to provide the most fruitful data based on information gathered by a network of low-cost sensors [20] spread about the region of consideration. Both systems are only able to observe local features. In our case study, the local features are images captured from downward-facing cameras on both of the platforms. The drifters used in our study have no control over the region in which they survey, but the ASV can make planar changes on the ocean’s surface to selectively sample. Our sensor nodes only report observations for the locally observed phenomena which they happen to float over, meaning that our estimate of the distribution is highly influenced by natural forces (wind, current, wave action) and

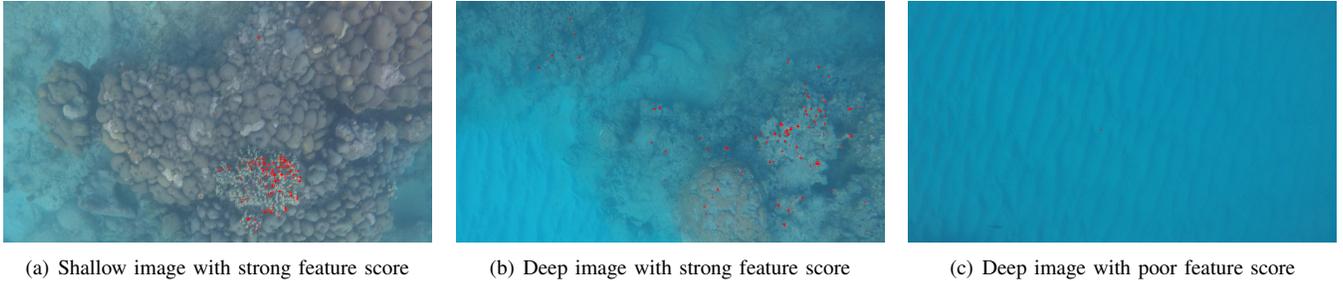


Figure 3: Examples of underwater images with ORB keypoints shown in red. These samples illustrate the benefit of selectively sampling a region so that we avoid areas with low information content. Feature score indicates how well we expect to be able to stitch overlapping images together to represent the ocean floor.

their starting positions.

Figure 1 presents an overview of our approach. In this figure, the drifter on the left (Drifter 1) observes an important feature in its local area and communicates this information to the surface vehicle. The drifter in the middle (Drifter 2) observes unimportant features and does not communicate any information to the ASV. The ASV then prioritizes these suggested locations according to a cost function which weighs relative distance and entropy and visits the most information-rich location. In this case the ASV will choose to travel to the left sensor (Drifter 1) location to closely sample, skipping the uninteresting areas, thus adaptively sampling the region of interest to maximize the knowledge gain with minimum distance traveled.

II. ASSIGNING A FEATURE SCORE TO VISUAL SAMPLES

We evaluate our adaptive sampling technique by setting up an experiment to strategically sample a region using visual features. The region considered has coral reef outcrops with an unknown distribution from which we would like to sample. Areas between corals are usually sand which is difficult to mosaic and uninformative. We wish to produce visually appealing surveys of sections (see Figure 8) of representative coral while minimizing the energy expenditure of our ASV which is carrying a quality camera needed for the survey. Our approach attempts to achieve this by prioritizing sampling in areas which are identified to have images that are assigned a high *feature score* by randomly dispersed drifters.

Underwater environments impose limitations on what can be captured with optical sensors and local water conditions greatly affect the quality of images captured. Perhaps the most basic reason that an image captured from the surface is of poor quality is because the scene is too deep or too shallow. This is a limitation of our surface observational capability that makes areas outside of some optimal depth range uninformative. If a depth map of the region is available, it can provide some insight into which areas may be effectively sampled, but there are other factors which influence the quality of the image that are difficult to determine

in advance. These include water clarity, natural illumination, and scattering which can make objects in an image difficult to ascertain [21]. Our experiments are conducted in relatively turbid water without the benefit of an artificial light source and thus many of these conditions are non-stationary over time and space.

In addition to quantifying how well we expect to see objects on the seafloor, we also want to determine if there is anything worth sampling at a particular point. Since our ultimate goal is to create appealing mosaics of coral outcrops (Figure 8), we use the number of keypoints detected in a sampled image as a metric of quality [22]. Image keypoints, which are pixels that represent distinctive points in an image, are matched across two images with overlapping scenes in feature-based mosaicking methods to calculate homography. Images captured over deep water, regions of solid sand, or in conditions with low visibility are often blurry or solid in colors and thus have few or no keypoints required to align them with their neighbors. On the other hand, observations over coral reef contain numerous distinctive corners, as seen in Figure 3. There are many well-known approaches to detecting keypoints in an image that vary widely in computational complexity and invariance. We chose to use an ORB (Oriented Fast and Rotated Brief) [23] feature detector because it is relatively lightweight and performs well in underwater scenes according to feature detector analysis presented by Quattrini Li *et al.* in [24].

In our system, each drifter captures images at a rate of $1Hz$ and then calculates the number of feature points in those images. The feature score is determined by dividing the number of features detected by a scaling factor such that its value is 1 for feature-rich images and 0 for images which fail to align with images of the same scene. Prior to deployment, we experimentally calculated a scaling factor for each camera to calibrate the system such that we can produce a feature score that is approximately equivalent across platforms. Figure 3 demonstrates several images captured in our experiment with relative depths and feature scores. If the feature score is over a certain threshold, the location

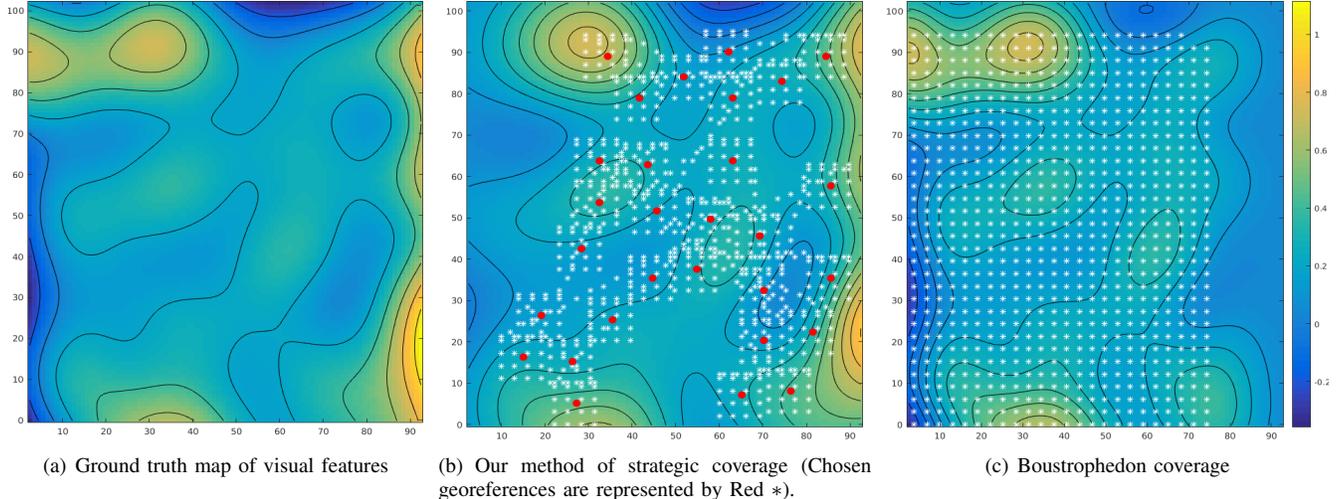


Figure 4: Ground truth and maps generated by extensive coverage (c) and our strategic sampling (b) after the surface vehicle has traveled 2400 m (sampling points are represented by White *). The axes are in meters.

is considered to be information-rich and drifter sends this feature score along with its georeferences to the ASV.

III. STRATEGIC SAMPLING BASED ON EXPECTED INFORMATION GAIN

Our approach seeks to intelligently sample a region of interest so as to reduce the time spent sampling areas which are unimportant for our goal. We make use of randomly dispersed, low-cost sensor nodes to improve our sophisticated vehicle’s understanding of the region. The vehicle uses this improved estimate of the underlying distribution to effectively plan a route so as to maximally sample information-rich areas. One of the significant features of our approach is that the communication between drifter and ASV happens asynchronously from planning and execution of the survey. Thus the plan can be adapted in real-time to improve the survey quality by including the recent location suggestions by the drifters.

Team of drifters start reporting feature scores and georeferences for regions which pass the feature score interestingness threshold. The ASV maintains a list, L , which contains all reported points and their respective scores. This list is continuously updated as information is received from drifters. The ASV chooses a point, l^* , to densely survey by finding a previously unvisited point in L which has the lowest associated cost as calculated by Eq. 1.

$$l^* = \arg \min_{l \in L} C_l \quad (1)$$

where the cost C_l used for waypoint ranking is given by,

$$C_l = d_l + \alpha(1 - s_l) \quad (2)$$

Eq. 2 weighs the relative distance, d_l , from the current position of the vehicle to location l with its reported feature

score s_l . The parameter α allows us to tune the weightage given to distance and the feature score at a given point. For instance, we can adaptively tune α based on the battery level of the vehicle, preferring nearby, but less interesting points when we are low on power. For our simulations we used $\alpha = 1$ as we weigh the feature score and the distance equally. But we would like to test the effect of this parameter in our future experiments.

Once the optimal point, l^* , is chosen, we execute a small, dense boustrophedon pattern which is centered over l^* , starting from the corner nearest to the ASV’s current position. This small survey helps in collecting continuous samples and not just one point-sample in the interesting area. After executing this dense coverage, the ASV will reconsider L according to its new position. While the ASV was executing its dense survey of the previous l^* , L is updated with new locations reported by the drifters and added by the ASV itself from the data collected during the transit.

IV. SIMULATIONS AND EXPERIMENTAL RESULTS

We evaluate our approach in simulation over ground-truth data gathered from field observations. This simulation consists of three main parts: the underlying feature distribution from which we sample, semi-synthetic drifter tracks, and the planning and execution of ASV paths. The underlying feature map (Figure 4(a)) is generated by densely sampling the survey site (see Section V for details on survey site) with a surface vehicle and building a map of feature scores using RBF-kerneled Gaussian Processes for interpolation.

Although we control the initial location from which a drifter begins its survey, we have little control over its path as it floats about the surface. To encapsulate realistic drifter movement in our simulation, we generated 43 synthetic

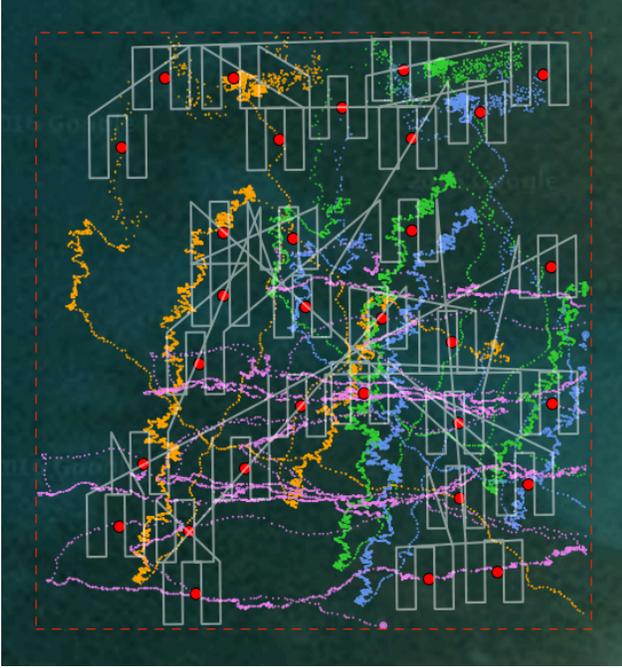


Figure 5: This figure represents the output from one of the simulation runs with four drifters (orange, lavender, blue, and green tracks indicate where a drifter captured an image). Red points indicate an area of interest that the ASV (tracks shown in white) chose to sample densely. The region of interest is represented by the bounding box with orange dotted lines ($100\text{ m} \times 90\text{ m}$).

drifter tracks by randomly shifting the start point of 5 real drifter tracks that were captured at the survey site. Thus, the simulated drifters exhibit realistic movement and timescales. Like a real drifter, a simulated drifter samples a feature score from the ground-truth distribution and if score passes the threshold for interestingness, it reports this information synchronously to the ASV. A threshold value of 0.25 was used in our simulations. The simulated ASV considers time and distance covered based on estimates from our real vehicle. Timestamps across the drifters and ASV are consistent with the realistic system, but sped up for experimental purposes. The ASV only considers drifter reports which were available prior to its current simulated time for consideration. For simplicity, we ignore non-holonomic constraints of the surface vehicle in simulation and also consider a complete communication between drifters and the ASV.

Figure 5 illustrates simulated trajectory paths for four drifters (depicted by orange, lavender, blue, and green points) along with the survey path planned and executed by the simulated robot (depicted in white lines). In this figure, the ASV has selected the red points to perform small, dense surveys. We evaluate our approach by comparing our

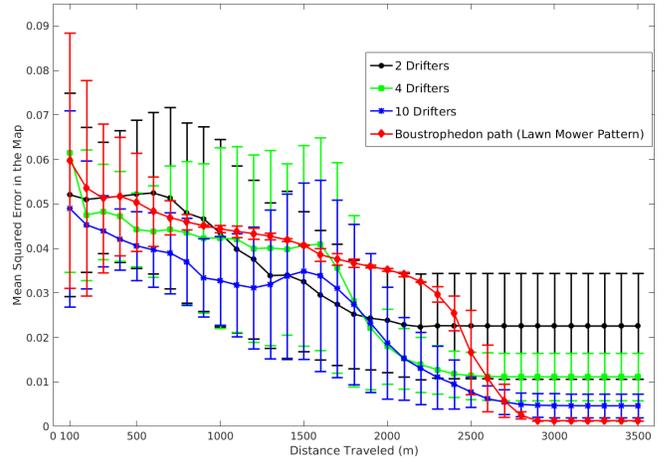


Figure 6: Plot of mean squared error in the map vs. distance traveled (m). We see that our model is able to predict values in the region quicker than an exhaustive search, though full coverage ultimately gives a clearer picture of the ground truth.

feature score map to the one generated by sampling on a traditional boustrophedon path. Figure 4(b) and Figure 4(c) are the maps generated by our approach and a traditional boustrophedon path after traveling a total of 2400 meters in simulation. The white markers indicate the sampling points. The chosen sampling points used in our algorithm is denoted by red markers in Figure 4(b). 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 ground-truth distribution (Figure 4(a)).

We evaluated the performance of our approach by running 100 simulations with a varying number of n randomly selected synthetic drifter tracks. For the considered survey area, n ranged from 2 to 10 with selected results depicted in Figure 6. The data in Figure 6 indicates the mean of the MSE of the feature score map generated after a distance traveled by the ASV compared against our ground truth. The solid lines represent the mean and error bars represent the standard deviation over 100 trials. We also consider 100 trials of boustrophedon paths starting at random corners and directions of the region. We exclude the first 10 meters of all surveys on this figure as the initial estimate is the same for all trials.

Boustrophedon paths cover a region one transect at a time, gradually collecting more information, regardless of the data, until the survey area is completely covered. We show this experimentally in the red line of Figure 6, which depicts the MSE for boustrophedon path slowly decreasing to zero as the ASV covers more area. In our approach, because the robot is driven to information rich locations

by the drifters, the initial few hundred meters of travel brings down the MSE significantly, but the system never achieves complete representation of the region. A careful observation of final MSE achieved with drifters reveals that more drifters tend to reduce the final MSE score. This makes sense, as more drifters will tend to observe more points in the region driving the vehicle to cover more terrain. Also due to uncontrollable motion of drifters, they will not visit all the locations in the region of interest. Thus with less number (≤ 4) of drifters deployed, some parts of the region remain unexplored resulting in non-zero MSE.

V. FIELD DEPLOYMENTS

We tested over a shallow region known to have several coral outcrops in the Folkestone Marine Reserve in Barbados. To build a baseline, we initially sampled the region extensively and built a feature score map which was utilized for generating the ground truth in our simulations.



Figure 7: Sampling points from our field experiment. Four drifter points (Black) were chosen by the ASV for dense sampling (Red).

To evaluate the strategic sampling approach presented in this paper, we placed four drifters (Figure 2(a)) randomly over the region and used a Clearpath Kingfisher (Figure 2(b)) as our ASV. Marine field deployments are always a challenge and are highly dependent on weather, current, and tidal conditions. Unfortunately, our field experiment was plagued by strong wind and low tides which stirred up sand and silt, greatly reducing visibility. This meant that our cameras were unable to see the reef during our trial and only reported feature scores of 0. This somewhat validates our work, as the ASV was not compelled to survey any of points over the unobservable reef.

Ultimately, we tested the algorithm by inducing strong feature scores by removing the drifters from the water so that they captured images above water. This prompted the ASV to correctly survey the areas around our artificial survey points as depicted in Figure. 7. In future field trials, we plan to perform additional experiments which test our approach

and expand our sensing capabilities so that we are less dependent on local meteorological conditions.

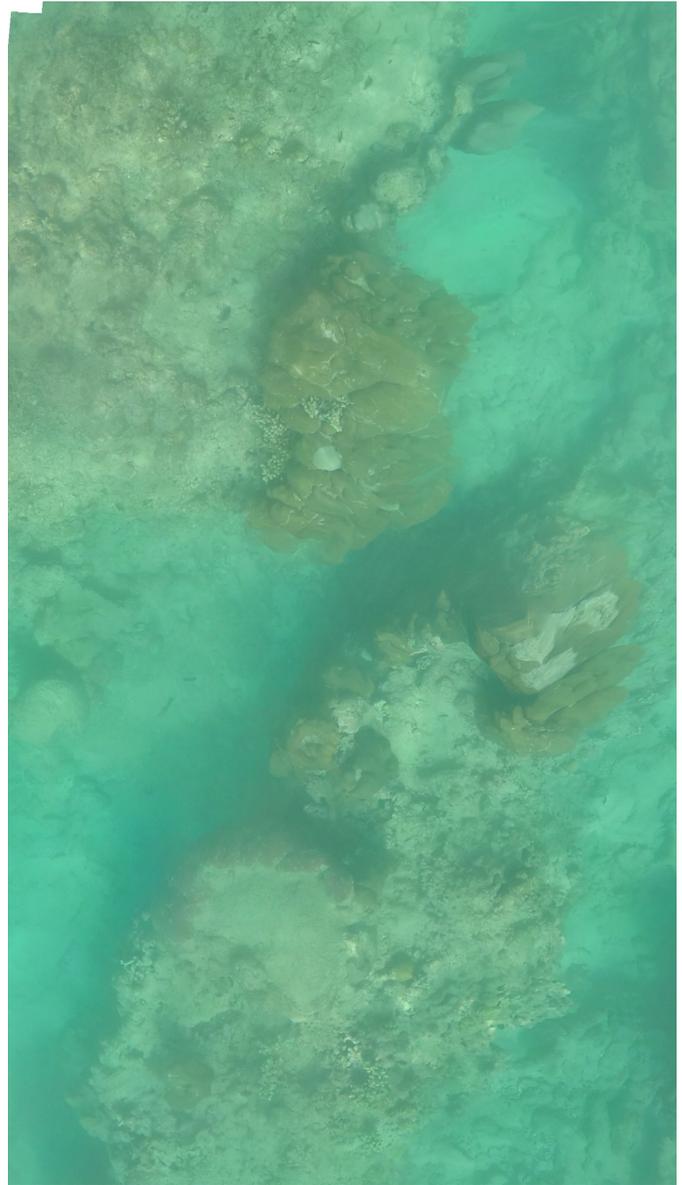


Figure 8: Mosaic developed using 27 images captured from ASV over coral heads. The mosaic is cropped to fit the column width.

VI. CONCLUSIONS

In this paper, we proposed a method for strategically sampling locally observable features from a spatially varying field using a tiered team of sensor platforms. This approach allows us to efficiently deploy an expensive, sophisticated sensor by deploying a team of low-cost sensors which provide insight that would otherwise be costly to capture. We evaluated our approach on a visual survey of a coral reef

and were able to effectively represent the underlying feature score map with less energy expenditure than exhaustive coverage. We are currently working towards developing hardware so that we can monitor other phenomena such as algae blooms or surface currents. In addition, we posit that this general approach will extend beyond marine applications. For example, we think it will extend nicely to systems which do not need to own the sensor node hardware. Consider a situation where a sophisticated wheeled or flying robot uses participating smart phones to passively provide information which is useful for achieving its goals.

In addition to extending potential applications, there are many improvements and research directions that are interesting to pursue. In this approach, we have assumed that communication is perfect, but this may not always be the case, especially when the survey area is large. Future work will consider routing the ASV so as to maximize the likelihood of achieving communication with nodes which have lost connection. We will also explore adding simple steering capabilities to the drifters so that they can have some control over their position, while still keeping cost down. After field experiments, we now understand that while our approach reduced the distance that the ASV needed to cover to effectively survey the area, the process of deploying and collecting disparate sensors requires considerably more human labor than an exhaustive pattern. Ideally the ASV could perform these tasks autonomously, though this will require significant additional hardware. On the algorithmic side, we will evaluate our approach against intelligent adaptive sampling strategies, which we did not have time to consider in this report. In addition, we recognize that the success of the proposed approach is largely dependent on the underlying feature distribution. Future effort will work to evaluate strategies for determining the most appropriate sampling behavior for an ASV as it learns the underlying distribution from sampling provided by the system. We future plan to incorporate additional planning strategies to optimize across additional sensor sources [16]. This will allow us to develop informed probabilistic paths to cover interesting points, rather than simply covering interesting points with small, boustrophedon patterns.

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 acknowledge funding for this work from the National Science Foundation (NSF 1513203, 1637876) and a Google Faculty Research Award. Our work would not have been possible without tremendous support from our colleagues. We are grateful for the wonderful staff at the Bellairs Research Institute whose helpfulness enabled us to conduct field experiments. We also appreciate our team members in the Mobile Robotics Lab at McGill University and Nare

Karapetyan in the Autonomous Field Robotics Lab at the University of South Carolina for technical discussions and assistance in field experiments.

REFERENCES

- [1] M. Xanthidis, A. Quattrini Li, and I. Rekleitis, "Shallow coral reef surveying by inexpensive drifters," in *MTS/IEEE OCEANS Shanghai*, China, Apr. 2016, pp. 1–9.
- [2] F. Shkurti, A. Xu, M. Meghjani, J. C. G. Higuera, Y. Girdhar, P. Giguere, B. B. Dey, J. Li, A. Kalmbach, C. Prahacs, K. Turgeon, I. Rekleitis, and G. Dudek, "Multi-domain monitoring of marine environments using a heterogeneous robot team," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Portugal, Oct. 2012, pp. 1447–1753.
- [3] S. B. Williams, O. Pizarro, J. M. Webser, R. J. Beaman, I. Mahon, M. Johnson-Roberson, and T. Bridge, "AUV-assisted Surveying of Drowned Reefs on the Shelf Edge of the Great Barrier Reef, Australia," *Journal of Field Robotics*, vol. 27, no. 5, pp. 675–697, 2010.
- [4] O. Pizarro, S. B. Williams, M. Johnson-Roberson, D. Steinberg, and M. Bryson, "Reef mapping and monitoring assisted by machines," in *Ocean Sciences 2014 Meeting*, 2014.
- [5] Y. Girdhar and G. Dudek, "Modeling curiosity in a mobile robot for long-term autonomous exploration and monitoring," *Autonomous Robots*, vol. 40, no. 7, pp. 1267–1278, 10 2016.
- [6] R. Brainard, R. Moffitt, M. Timmers, G. Paulay, L. Plaisance, N. Knowlton, J. Caley, F. Fohrer, A. Charette, C. Meyer *et al.*, "Autonomous reef monitoring structures (arms): A tool for monitoring indices of biodiversity in the pacific islands," in *11th Pacific Science Inter-Congress, Papeete, Tahiti*, 2009.
- [7] M. Meghjani, S. Manjanna, and G. Dudek, "Multi-target rendezvous search," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016. IEEE, 2016, pp. 2596–2603.
- [8] I. T. MacDonald and J. C. Mullarney, "A novel floccdrifter platform for observing flocculation and turbulence processes in a lagrangian frame of reference," *Journal of Atmospheric and Oceanic Technology*, vol. 32, no. 3, pp. 547–561, 2015.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] H. Choset and P. Pignon, “Coverage path planning: The boustrophedon cellular decomposition,” in *In International Conference on Field and Service Robotics*, 1997.
- [13] 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.
- [14] 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.
- [15] 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.
- [16] S. Manjanna, N. Kakodkar, M. Meghjani, and G. Dudek, “Efficient terrain driven coral coverage using gaussian processes for mosaic synthesis,” in *13th Conference on Computer and Robot Vision (CRV), 2016*. IEEE, 2016, pp. 448–455.
- [17] M. Meghjani, F. Shkurti, J. C. G. Higuera, A. Kalmbach, D. Whitney, and G. Dudek, “Asymmetric rendezvous search at sea,” in *Canadian Conference on Computer and Robot Vision (CRV)*. IEEE, 2014, pp. 175–180.
- [18] 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,” *IEEE Journal on Selected Areas in Communications*, vol. 30, no. 5, pp. 899–911, 2012.
- [19] Z. Ma, K. Yin, L. Liu, and G. S. Sukhatme, “A spatio-temporal representation for the orienteering problem with time-varying profits,” *arXiv preprint arXiv:1611.08037*, 2016.
- [20] D. Boydston, M. Farich, J. M. III, S. Rubinson, Z. Smith, and I. Rekleitis, “Drifter sensor network for environmental monitoring,” in *12th Conference on Computer Robot Vision*, Halifax, Nova Scotia, Canada, June 2015, pp. 16–22.
- [21] O. Pizarro and H. Singh, “Toward large-area mosaicing for underwater scientific applications,” *IEEE Journal of Oceanic Engineering*, vol. 28, no. 4, pp. 651–672, Oct 2003.
- [22] D. Bekele, M. Teutsch, and T. Schuchert, “Evaluation of binary keypoint descriptors,” in *2013 IEEE International Conference on Image Processing*, Sept 2013, pp. 3652–3656.
- [23] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, “Orb: An efficient alternative to sift or surf,” in *Proceedings of the 2011 International Conference on Computer Vision*, ser. ICCV ’11. Washington, DC, USA: IEEE Computer Society, 2011, pp. 2564–2571. [Online]. Available: <http://dx.doi.org/10.1109/ICCV.2011.6126544>
- [24] A. Quattrini Li, A. Coskun, S. M. Doherty, S. Ghasemlou, A. S. Jagtap, M. Modasshir, S. Rahman, A. Singh, M. Xanthidis, J. M. OKane, and I. Rekleitis, “Vision-based shipwreck mapping: on evaluating features quality and open source state estimation packages,” in *MTS/IEEE OCEANS Monterey*, Sept. 2016, pp. 1–10.