Online Summaries as a Framework for Perception and Planning in Marine Robotic Systems

Jeff Kaeli, Gwyneth Packard, Matt Grund, Carl Fiester, Judson Poole, Dara Tebo, Mike Purcell
Ocean Systems Laboratory, Applied Ocean Physics and Engineering Department
Woods Hole Oceanographic Institution (WHOI)
Woods Hole, MA, USA
{jkaeli, gepackard, mgrund, cfiester, jpoole, dtebo, mpurcell}@whoi.edu

***Abstract*—We present a scalable framework that enhances the perception and planning capabilities of Autonomous Underwater Vehicles (AUVs). This framework has two key attributes. First, we use online summaries as a real-time data reduction tool for maintaining a current representation of the data collected at any point during a mission. Second, the data points from this summary are stored in a database and indexed by their spatial and temporal metadata so they can be queried for planning. We demonstrate how this simple yet powerful approach can reduce the amount of data required to represent a data set to a specified resolution.**

***Index Terms*—autonomous underwater vehicle, autonomy, perception, data management**

I. Introduction

Autonomous Underwater Vehicles (AUVs) are revolutionizing our ability to sample and understand the ocean. Their "autonomy" comes from their programmed ability to react to data as it is collected in real time during missions. For the simplest missions, an AUV might follow preprogrammed track lines to survey an area of interest, using navigation sensors to stay on course and maintain an appropriate depth in the water column. Two active areas of autonomy research aim at developing algorithms that give the AUV some freedom to decide what path to take to successfully execute its objectives. The first is focused on path planning and navigation based on some known map or environmental model of an area. The second is focused on adaptive behaviors that are a reaction to some event that has occurred in its data streams.

A notable characteristic of these approaches is that the AUV is acting upon either a static model or a single observation in space and time. There is no mechanism to update this representation with new information, nor is there a robust method for tying multiple events together in a unified representation of the world. Our work attempts to bridge the gap between these two approaches using an online summary framework. Traditional clustering algorithms are "offline" summaries, such as the popular K-means algorithm, meaning that they operate on entire data sets that have already been collected. In contrast, an "online" summary is a clustering technique that operates on data points one-at-a-time and in the order they are presented to the algorithm. This makes them particularly well suited to applications in robotics because they enable the robot to maintain a dynamic "summary" of it environment based on its observations up to the current time.

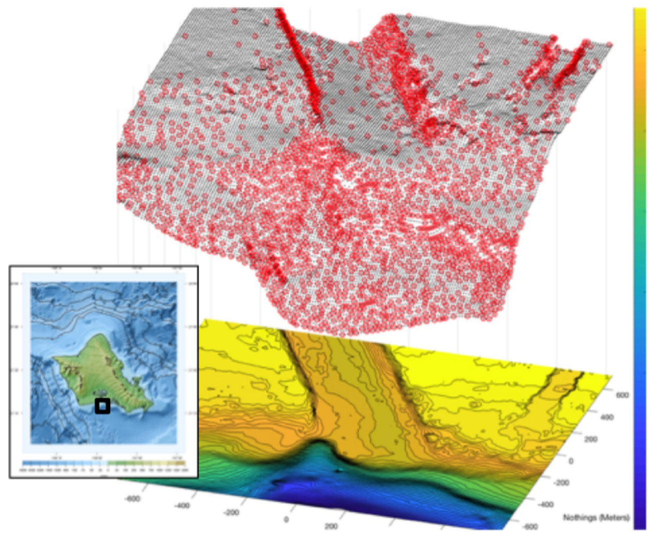

Fig. 1. Summarizing raster data using a set of soundings. [1]

Our representation can be thought of as a set of soundings that serves both as a data reduction tool as well as a tool for constructing locally constrained queries useful for in situ mission planning, making it fast, lightweight, and scalable to many kinds of data. Figure 1 shows a gridded bathymetry data set that has been summarized by our algorithm using a set of soundings. In this paper, We first present the mathematical formulations governing the online summary algorithm, building on prior work in the literature. Next, we describe the database schema employed to store the summary data points and their locations in space and time. We then go on to discuss our implementation of this framework on an AUV and present results from open water missions using bathymetry data. We discuss how the same framework can be applied to data offline as either a pre-mission or a post-mission analysis tool.

II. Related Work

A. Path Planning and Navigation

Path planning is a well-studied problem in terrestrial robotics [2]. In the underwater domain, autonomous path planning is largely motivated by a desire to avoid obstacles such as shallow water or to optimize for some environmental variable such as currents [3]. One approach is to first perform some sort of maneuver to establish an environmental baseline, in this case current direction, and then adapt the survey track lines to run parallel to the current direction [4]. Because of the challenges associated with localization underwater, using a priori maps as a navigation aid has been studied [5]. Many of these terrain relative navigation approaches have origins in missile guidance and space craft landing applications [6]. However, these approaches rely on a prior maps, which can be costly to create. Simultaneous Localization and Mapping (SLAM) algorithms have been applied underwater to both constrain navigation uncertainty while building or augmenting a map of an unfamiliar environment [7]. However, SLAM techniques perform poorly in environments lacking rich features.

B. Adaptive Behaviors

Many adaptive AUV behaviors are pre-scripted maneuvers triggered by detecting some feature in the environment [8]. Another application of adaptive AUV behavior is characterizing and tracking dynamic physical oceanographic features such as fronts [9] and plumes [10]. In one case, a reactive behavior is triggered by measured increase in propane concentration [11]. Such capabilities are very important as the Arctic because increasingly ice-free, motivating the need for long range AUVs capable of detecting and monitoring oil spills [12]. However, the majority of these techniques still operate based on data that is in the current field of view of the sensor, leaving large-scale visualization and understanding to post-processing analysis.

C. Mission Summaries

Because of the large volume of data collected by autonomous platforms, there is a strong motivation to summarize that data, either for operator understanding or to inform adaptive vehicle behaviors. Image-based summaries for robots have been demonstrated using topic modeling [13] as well as using the concept of "surprise" as a metric for information gain [14]. Bayesian surprise, as quantified by the Kullback-Leibler divergence, has been shown to correlate well with human attention [15]. This has been applied to AUV surveys for the purposes of semantic data set compression as well as anomaly detection [16]. While these approaches aim to summarize data in feature space, the current effort builds upon this work by extending it to summaries in physical space to generate sparse, multi-resolution maps.

III. Online Summaries

A. Formulation

Each data point collected by the robot can be thought of as an observation $\mathbf{X}$ containing a measurement vector $\mathbf{M}$ (i.e. sensor data) of dimension G and a location vector $\mathbf{L}$ (i.e. spatiotemporal metadata) of dimension H.

$$\mathbf{M} = [m_1, m_2, ...m_G] \tag{1}$$

$$\mathbf{L} = [l_1, l_2, ...l_H] \tag{2}$$

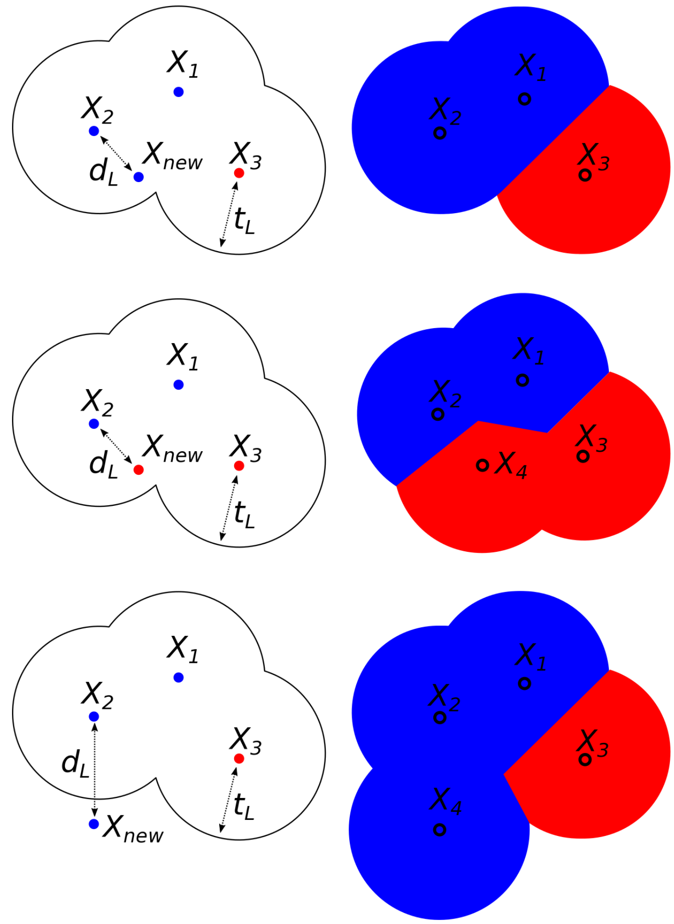

Fig. 2. Cartoon depicting the online summary algorithm for a binary measurement (**red** or **blue**) with a two-dimensional location space. Each row shows a new observation relative to summary set (left) and the resulting model of the world after an exhaustive K=1 nearest-neighbor query (right). *Top:* The new observation is **blue** which matches the nearest member of the summary set. The current summary set accurately predicts the new observation, so the new observation is not added to the summary set. *Center:* The new observation is **red** which does not match the nearest member of the summary set. The current summary set does not accurately predict the new observation, so the new observation is added to the summary set. *Bottom:* The new observation is **blue** but this is irrelevant because the location distance to the nearest member of the summary exceeds the location distance threshold. The current summary set does not predict beyond this distance, so the new observation is automatically added to the summary set.

$$\mathbf{X} = \{\mathbf{M}, \mathbf{L}\} \tag{3}$$

The summary comprises a set $\mathbf{S}$ of N discrete observations or "soundings" that represent the robots knowledge of the world. These soundings can be interpolated to determine the predicted measurement value at any location.

$$\mathbf{S} = \{\mathbf{X}_1, \mathbf{X}_2, ...\mathbf{X}_N\} \tag{4}$$

In the G-dimensional space of possible measurements, and in the H-dimensional space of possible locations, we can define distance metrics D_M and D_L that return a distance d_M or d_L between two measurements or two locations, respectively.

$$d_{M,1-2} = D_M(\mathbf{M}_1, \mathbf{M}_2) \tag{5}$$

$$d_{L,1-2} = D_L(\mathbf{L}_1, \mathbf{L}_2) \tag{6}$$

We can also define a query function Q_S for the summary set that returns an interpolated measurement at a given location. An important parameter of this function is K, the number of soundings required to compute the interpolation.

$$\mathbf{M}_{interp} = Q_S(\mathbf{L}_{query}|K) \tag{7}$$

When a new observation $\mathbf{X}_{new} = \{\mathbf{M}_{new}, \mathbf{L}_{new}\}$ is made, the summary set is queried at the new location. The distance is computed between the newly observed measurement and the interpolated measurement.

$$d_{M,new-interp} = D_M(\mathbf{M}_{new}, Q_S(\mathbf{L}_{new}|K)) \tag{8}$$

If this distance is less than or equal to a specified measurement threshold t_M, then the new observation can be said to be already characterized by the existing summary set. However, if this distance exceeds this threshold, then the new observation offers novel information and should be added as a new member of the summary set. A location distance threshold t_L can also be set such that any new observations whose location distance d_L among the K nearest summary set members exceeds this threshold is automatically added to the summary set. Figure 2 illustrated several examples of this process. In this way, thresholds are analogous to resolutions in the measurement and location spaces.

B. *Implementation*

The algorithm was developed in MATLAB and then implemented in C++ for efficient real-time execution. The summary set of soundings is stored in memory as an unsorted array. The computational costs of searches of this array are linear in the size of the summary set N. While this suffers in performance as the summary set grows, there is ultimately a performance gain because the array does not have to be re-indexed when each novel observation is added. Future implementations could utilize a buffer to store $N_B \leq N$ unsorted new soundings alongside an existing sorted summary set. Queries will then be at worst $N_B + \log(N)$ in complexity. When this buffer is full, the summary can be re-indexed with $2N$ soundings and the buffer cleared.

Bathymetry was selected as an initial development case because it has direct relevance to adaptive path planning, it is a ubiquitous measurement made by AUVs (depth plus altitude), and it is simple to model as a 1-D measurement $\mathbf{M} = z$ at a 2-D location $\mathbf{L} = [x, y]$. The distance metric in measurement space is the absolute value of the difference between two bathymetry values.

$$d_{M,1-2} = |z_1 - z_2| \tag{9}$$

Fig. 3. The REMUS 600 AUV.

The distance metric in location space is the Cartesian distance between two locations.

$$d_{L,1-2} = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2} \tag{10}$$

We experimented with several query and interpolation functions including the simple K=1 nearest neighbor. The results presented here use a K=3 nearest neighbor with an inverse-distance weighting scheme such that the limit as a queried location approaches a sounding location equals the measurement at that sounding, and the limit as distance increases approaches the mean of the K measurements.

IV. Results

A. *Local Testing*

Local testing was conducted in Vineyard Sound aboard the R/V Discovery using a REMUS 600 AUV shown in Figure 3 owned and operated by the Woods Hole Oceanographic Institution (WHOI). This vehicle has a new core electronics board designed for power efficiency and a PHINS inertial navigation system [17]. The bathymetry of the survey area is gently sloping to the southeast between roughly 5 and 10 meters of depth. Two missions were programmed onto the vehicle, shown in Figure 4 at left. Mission 1 was a spiral survey pattern with track line spacing of 30 meters between revolutions. Mission 2 was a traditional mow-the-lawn type mission with track spacing of 50 meters and legs in first the north-south and then the east-west orientations.

Depth and altitude values were obtained from a pressure sensor and an ADCP/DVL, respectively. The values were added to obtain bathymetry and then time averaged over 20 vehicle cycles or approximately 2 seconds. This reduced the processor load while naively removing some noise and outliers from the data stream. As a rudimentary quality control step, minimum and maximum allowable values for bathymetry were set at 3 and 60 meters, respectively. The measurement threshold t_M was set to 0.5 meters, while the location threshold was set to 50 meters. Mission 1 began with an empty summary set, analogous to having no prior information of the survey area. Mission 2 began with the summary set generated from Mission

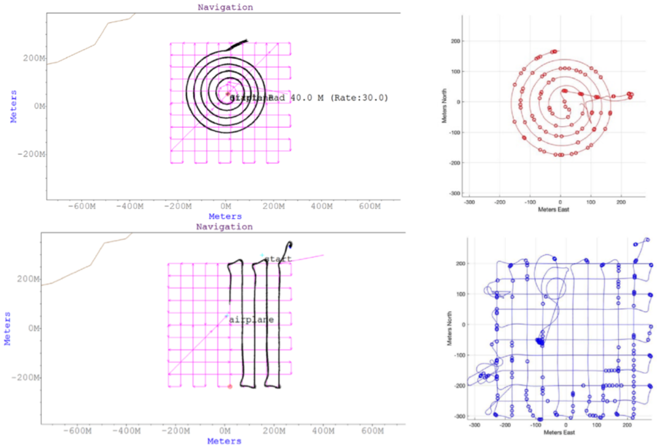

Fig. 4. The planned track lines (*left*) and saved soundings (*right*) from Mission 1 (*top*) and Mission 2 (*bottom*).

1, analogous to carrying an a partial a priori map of the survey area. As is shown in the right half of Figure 4, the first mission produced a relatively uniform set of 90 soundings, and the second mission added a similarly uniform 191 soundings, primarily in areas that had not yet been surveyed. During this second mission, the vehicle has access to the prior knowledge of the world while simultaneously updating that corpus of knowledge with novel information.

Although the algorithm simply saves single bathymetry measurements to the summary set, the summary can be queried at any time on a grid to form a raster image. This is useful both for planning algorithms that operate on regular grids as well as for human visualization of the data. Figure 5 shows a progression of these raster maps as they would appear to the vehicle at various points throughout each mission. As more information is gathered and more soundings are saved to the summary set, a more complete picture of the environment is formed.

The location threshold t_L is a crucial parameter, both when building a summary in real time as well as when querying the summary to build a raster map. In the first case, new soundings will be added whenever the distances to the K-nearest soundings lie beyond this threshold. The summary set created during Missions 1 and 2 would be sparser if the location threshold was set to a higher value, largely because, in the case of the gently sloping seafloor in this region, bathymetry values could be predicted from soundings that are further away. In the second case, creating a raster map with a location threshold that is too small will result in gaps in the map. If the location threshold is set to infinity, then any queried location arbitrarily far away will return some "best guess" of the map value as informed by the interpolation function. This may either be desirable or dangerous, depending on the application. For instance, an on-board planner might prefer to receive an undefined value in opposed to a highly uncertain extrapolated value.

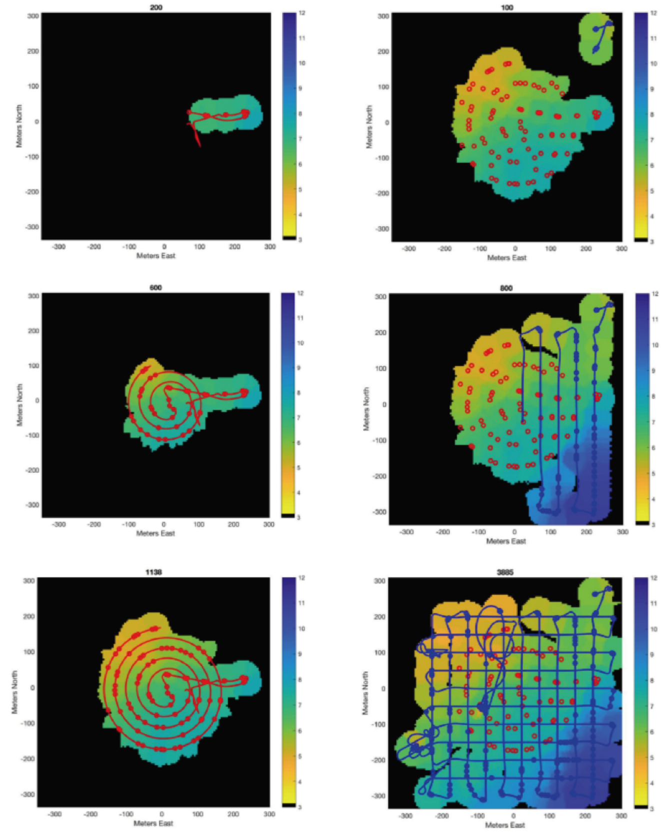

Fig. 5. The queried raster map as it is progressively built through Mission 1 (red, left) and Mission 2 (blue, right). Circles denote locations have been added to the summary set.

V. Discussion

In this framework, new soundings are only added to the summary set when they contain novel information. This novelty manifests itself as a threshold exceeded in either the location space or in the measurement space. In a manner of speaking, this is a form of anomaly detection, where the anomalies represent observations that differ from the current knowledge of the world. We can look at the difference between two summary sets and compare them to see what anomalies exist, or how the understanding of the world has changed between the two. Figure 6 illustrates this concept, showing the absolute difference between the raster map generated at the completion of Mission 1 and the raster map generated at the completion of Mission 2.

While much of the difference map falls within the specified 0.5-meter measurement threshold, there are two notable exceptions. First, there are large errors at the edges of the map. This can be explained using the metaphor of extrapolating data at a cliff. After Mission 1, the best guess of these bathymetry values is that they are similar to those that have already been seen in the interior of the map. During Mission 2, these regions are actually explored, and their values are deeper than those predicted. One way to mitigate this is to redefine the query function so it better handles extrapolating data. The other

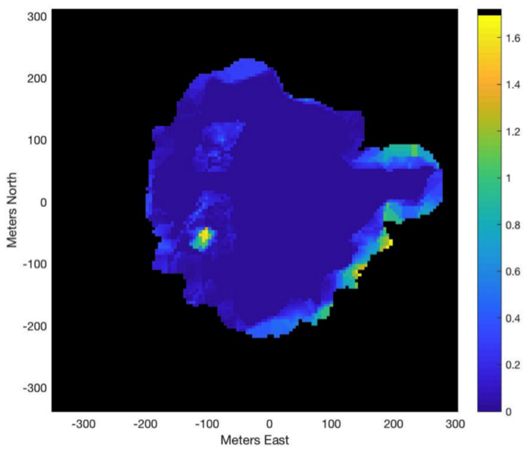

Fig. 6. Absolute differences between densely queried bathymetry maps generated from Mission 1 and Mission 1.

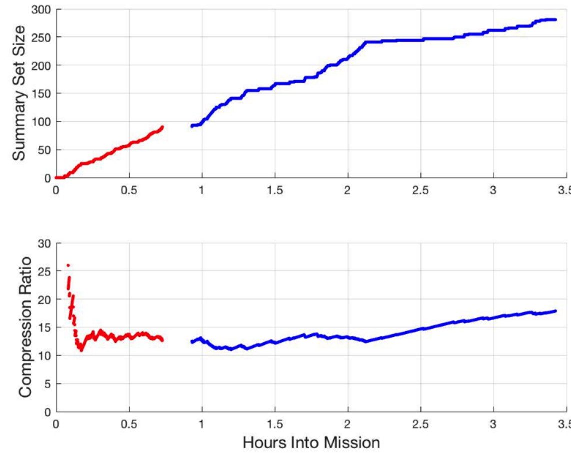

Fig. 7. Summary set size (*top*) and corresponding compression ratios (*bottom*) for Mission 1 (*red*) and Mission 2 (*blue*).

notable anomaly is a bright spot in the lower left of the survey area. Has a large rock appeared between Missions 1 and 2? Unlikely. What happened resulted from water ingress into a prototype 3D printed nose cone, affecting the pitch stability and navigation estimate of the vehicle. This offers insight into the necessity for good quality control of the data and highlights the importance of vehicle self-monitoring. It also illustrates the difficulty of diagnosing poor performance from a subset of data streams.

Furthermore, the issue of accurate navigation is a concern. The algorithm might confuse a non-novel observation at an incorrect location as a novel observation. State-of-the-art inertial navigation systems are typically accurate to around .1% of distance traveled, or one meter every kilometer [18]. This is sufficient for some applications, but will fail when the observation resolution is less than the distance threshold. One approach to this problem is to make locally consistent maps over short time scales and then fuse those sub-maps together [19]. We hope to address the problems in follow-on efforts by representing uncertainty alongside the observations in the summary set.

The underwater environment places extreme restrictions on the available bandwidth for communication. One of the attractive qualities of online summaries is its ability to reduce the file size required to represent a data set to a given resolution. Figure 7 shows the size of the summary set over Missions 1 and 2, as well as the effective compression ratio achieved but this summary. The ratio is artificially to start high because the quality control filter rejected the first several data points early in the mission. The ratio stabilizes as the summary set grows linearly as a result of the vehicle consistently exploring new terrain that is fairly uniform. When the summary set growth levels off, such as during Mission 2 when the vehicle is resurveying the previously mapped area, the compression ratio increases. The distance thresholds understandably have a direct effect on this compression ratio, as a map with coarser resolution can be stored using fewer bits.

Another useful feature of this framework is that the same basic mechanism can be used to generate a summary set from a priori information of an area. This was illustrated in Figure 1, where a raster bathymetry data set [1] with interesting bathymetric features has been fed through the algorithm to generate a set of soundings. The measurement threshold was set to 1 meter while the location threshold was set to infinity. Notably, the soundings are denser where the variation in bathymetry is more dramatic, and sparser in flat areas.

This exercise highlights an important characteristic of online summaries. After a single pass of the algorithm over the data set, a second pass might continue to add soundings to the summary set. This is because the order in which observations are presented to the algorithm matters, and while on the first pass an observation might fall within the threshold, a nearby data point added at a later time may cause the interpolated value of that observation to fall outside the threshold. This "jitter" in the map-making process is unsettling to us because we are accustomed to viewing whole data sets after they are collected when we have access to numerous tools for "optimal" analysis.

Online methods have less guarantees about optimal performance and are highly dependent on the data they are summarizing. We can, in theory, always devise an a priori map that would force the online summary to consider every new observation as novel regardless of the thresholds. Thankfully, however, data encountered in the real world is highly redundant, so with realistically tuned parameters massive shifts in the map are unlikely. In the case illustrated in Figure 1, the summary converged in six iterations, using a mere 2,777 soundings to summarize 21,291 gridded bathymetry points to a resolution of 1 meter.

These compression rates could further be improved upon by parameterizing the thresholds based on location. If navigation and obstacle avoidance are the ultimate goals, then smaller thresholds and higher resolutions are more important along shallow coastlines, and they can be more relaxed in deeper waters where there is more space to maneuver. Another application is understanding ocean structures such as plumes, fronts, or currents. In this case, the measurement space could be multidimensional (temperature, salinity, multiple velocity components) and the location space could be four dimensional including depth and time. We hope that future extensions to this work will couple summaries in the physical location space with existing work in feature-space summaries to more compactly represent complex structures in the environment.

ACKNOWLEDGMENT

We would like to thank the crew of the R/V Discovery for their invaluable help in conducting field tests and deploying and recovering the AUV.

REFERENCES

[1] "Oahu, hawaii 1/3 arc-second mhw coastal digital elevation model," www.ngdc.noaa.gov/mgg/dat/dems/regional_tr, 2011.

[2] L. Yang, J. Qi, D. Song, J. Xiao, J. Han, and Y. Xia, "Survey of robot 3d path planning algorithms," *Journal of Control Science and Engineering*, vol. 2016, p. 5, 2016.

[3] C. Petres, Y. Pailhas, P. Patron, Y. Petillot, J. Evans, and D. Lane, "Path planning for autonomous underwater vehicles," *IEEE Transactions on Robotics*, vol. 23, no. 2, pp. 331–341, 2007.

[4] D. P. Williams, F. Baralli, M. Micheli, and S. Vasoli, "Adaptive underwater sonar surveys in the presence of strong currents," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 2604–2611.

[5] D. K. Meduna, S. M. Rock, and R. McEwen, "Low-cost terrain relative navigation for long-range auvs," in *OCEANS 2008*. IEEE, 2008, pp. 1–7.

[6] A. E. Johnson and J. F. Montgomery, "Overview of terrain relative navigation approaches for precise lunar landing," in *2008 IEEE Aerospace Conference*. IEEE, 2008, pp. 1–10.

[7] J. J. Leonard and A. Bahr, "Autonomous underwater vehicle navigation," in *Springer Handbook of Ocean Engineering*. Springer, 2016, pp. 341–358.

[8] L. Freitag, M. Grund, C. Von Alt, R. Stokey, and T. Austin, "A shallow water acoustic network for mine countermeasures operations with autonomous underwater vehicles," *Underwater Defense Technology (UDT)*, pp. 1–6, 2005.

[9] S. M. Petillo and H. Schmidt, "Autonomous and adaptive underwater plume detection and tracking with auvs: Concepts, methods, and available technology," 2016.

[10] J. A. Farrell, S. Pang, and W. Li, "Chemical plume tracing via an autonomous underwater vehicle," *Oceanic Engineering, IEEE Journal of*, vol. 30, no. 2, pp. 428–442, 2005.

[11] R. Camilli, B. Bingham, M. Jakuba, H. Singh, and J. Whelan, "Integrating in-situ chemical sampling with auv control systems," in *Oceans' 04 MTS/IEEE Techno-Ocean'04 (IEEE Cat. No. 04CH37600)*, vol. 1. IEEE, 2004, pp. 101–109.

[12] A. Kukulya, J. Bellingham, J. Kaeli, C. Reddy, M. Godin, and R. Conmy, "Development of a propeller driven long range autonomous underwater vehicle (lrauv) for under-ice mapping of oil spills and environmental hazards: An arctic domain center of awareness project (adac)," in *2016 IEEE/OES Autonomous Underwater Vehicles (AUV)*. IEEE, 2016, pp. 95–100.

[13] R. Paul, D. Rus, and P. Newman, "How was your day? online visual workspace summaries using incremental clustering in topic space," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 4058–4065.

[14] Y. Girdhar and G. Dudek, "Efficient on-line data summarization using extremum summaries," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 2012, pp. 3490–3496.

[15] L. Itti and P. F. Baldi, "Bayesian surprise attracts human attention," in *Advances in neural information processing systems*, 2005, pp. 547–554.

[16] J. W. Kaeli and H. Singh, "Online data summaries for semantic mapping and anomaly detection with autonomous underwater vehicles," in *OCEANS 2015-Genova*. IEEE, 2015, pp. 1–7.

[17] C. Fiester, D. Gomes-Ibanez, M. Grund, M. Purcell, F. Jaffre, N. Forrester, T. Austin, and R. Stokey, "A modular, compact, and efficient next generation remus 600 auv," in *OCEANS*. MTS/IEEE, 2019, pp. 1–6.

[18] L. Paull, S. Saeedi, M. Seto, and H. Li, "Auv navigation and localization: A review," *IEEE Journal of Oceanic Engineering*, vol. 39, no. 1, pp. 131–149, 2013.

[19] C. Roman and H. Singh, "A self-consistent bathymetric mapping algorithm," *Journal of Field Robotics*, vol. 24, no. 1-2, pp. 23–50, 2007.