An Intelligent Approach to Hysteresis Compensation while Sampling using a Fleet of Autonomous Watercraft

Abhinav Valada, Christopher Tomaszewski, Balajee Kannan,
Prasanna Velagapudi, George Kantor, and Paul Scerri

The Robotics Institute, Carnegie Mellon University
Pittsburgh, PA 15213, USA

Abstract. This paper addresses the problem of using a fleet of autonomous watercraft to create models of various water quality parameters in complex environments using intelligent sampling algorithms. Maps depicting the spatial variation of these parameters can help researchers understand how certain ecological processes work and in turn help reduce the negative impact of human activities on the environment. In our domain of interest, it is infeasible to exhaustively sample the field to obtain statistically significant results. This problem is pertinent to autonomous water sampling where hysteresis in sensors causes delay in obtaining accurate measurements across a large field. In this paper, we present several different approaches to sampling with cooperative vehicles to quickly build accurate models of the environment. In addition, we describe a novel filter and a specialized planner that uses the gradient of sensor measurements to compensate for hysteresis while ensuring a fast sampling process. We validate the algorithms using results from both simulation and field experiments with four autonomous airboats measuring temperature and dissolved oxygen in a lake.

Keywords: Adaptive sampling, Active learning, Multi-robot systems, Autonomous surface vehicle, Environmental monitoring

1 Introduction

Recent advances in Autonomous Surface Vehicle (ASV) technology have enabled these systems to be used in missions that involve sampling large bodies of water for extended periods in order to monitor dynamic spatial and temporal phenomena with little or no human supervision. Monitoring water bodies is not only important for understanding the physiology of aquatic life but also for understanding how these systems are affected by both natural changes in the environment such as storms and volcanic eruptions as well as human activities such as surface run off from farms and industrial discharges. By collecting spatially distributed samples and analyzing the data it may be possible to predict how some of these processes work and potentially prevent adverse ecological effects such as eutrophication, oxygen depletion, and accelerated aging. ASVs are a natural choice for this kind of application as they have the capability to sample large areas while providing real-time measurements. They have been successfully used for mapping applications both above and below the water surface [1], even

at varying depths [2]. Cooperative fleets of ASVs have advantages over a single ASV in reliability, coverage and fault tolerance. Moreover intelligent sampling techniques can greatly improve the efficiency and quality of sampling by adaptively determining the next sampling locations based on the previously measured data.

A phenomenon known as hypoxia occurs in water bodies when oxygen saturation falls within the range of 1% to 30%, in which most types of fish and some invertebrates die due to insufficient oxygen. In the aquaculture industry, overstocking of fish increases susceptibility to hypoxia. Every year, this leads to millions of dollars worth of fish losses. The Fish and Wildlife Conservation Commission has recorded about 5114 incidents of fish kills from 1972 to 2012, where the number of fish lost at each incident ranged from hundreds to millions. There are several factors that lead to hypoxia, the most common being algae blooms, pollution, red tide, and rapid fluctuations in temperature. Water temperature in particular influences several aquatic processes such as the metabolic rates of organisms, level of dissolved oxygen and rate of photosynthesis, hence temperature fluctuations can not only drastically affect dissolved oxygen, but also cause severe imbalance in the aquatic ecosystem. Periodic monitoring of dissolved oxygen and temperature in aquacultural applications can help prevent such detrimental effects, as aerators can be deployed to replenish oxygen and heaters can be used to maintain the temperature. However, observed spatial and temporal variation patterns of dissolved oxygen and temperature are complex: they vary with the amount of dissolved solids, salinity and hydrodynamics of the water body, changing over the course of even a single day. Thus, these parameters must be sampled frequently in order to build accurate models.

In this paper, we describe our work on developing algorithms to adaptively sample dissolved oxygen and temperature using a fleet of autonomous watercraft. Initial experiments with autonomous sampling revealed that slow response of the polyethylene membrane used in dissolved oxygen sensors causes rate-dependent hysteresis, which significantly affects measurement accuracy. Similarly, with temperature sensors hysteresis is often caused by the introduction of some amount of strain or moisture penetrating inside the sensor. As these sensors are not specifically designed for dynamic measurements, a lag in the response of the sensor causes erroneous measurements if the vehicle travels at a rate that does not allow the sensor readings to stabilize. Developing a sensor model that predicts the rate of change is one potential solution, but even small errors in this estimate can dramatically affect the final measurement value. In our work we adopt an approach in which we use the time derivative of measurements, rather than the measurements themselves, to alter a pre-defined set of bounds that converge to the true value over time. The vehicle then plans its path based on the expected sensor value, range of the bounds, and the frequency with which the cell has been previously visited. To validate this approach we compare it to a suite of other sampling algorithms including random walks, lawn-mower patterns, level sets and maximum uncertainty sampling.

The rest of the paper is organized as follows. In Section 2 we formulate the problem in question and identify the associated related work in Section 3. We outline our bounding filter solution for modeling the environment in Section 5 and apply it to the problem of multi-robot planning in Section 6. In Section 7 we outline comparative algorithms that we implemented for comparison against our bounded filter approach and subsequently analyze the performance results from

the conducted experiments in Section 8. Finally, in Section 9, we conclude the paper with a summary of the results and an outline of future work directions.

2 Problem Formulation

In this work, we are interested in measuring a field which is continuous over $\mathbb{R}^n$. Thus we can define a field mapping function:

$$\Psi : \mathbb{R}^n \rightarrow \mathbb{R}$$

We model a sensor with hysteresis as a process with internal state, based on a simple exponential average:

$$s(t+1) = \alpha^t \cdot \Psi(x) + (1 - \alpha^t) \cdot s(t)$$

In continuous time, this corresponds to a simple first-order differential process:

$$\frac{\partial s}{\partial t} = \log(\alpha) (\Psi(x) - s(t))$$

Our objective is to estimate the value of $\Psi(x)$ at some set of points, $M \subseteq \mathbb{R}^n$. A simple method for doing so is simply to visit each point $x \in M$, and wait a certain amount of time for the sensor output s to approach $\Psi(x)$. However, α is small, $\frac{\partial s}{\partial t}$ is also small, meaning it can take a long time for s to approach $\Psi(x)$.

Additionally, there is a time cost associated with moving between the points in M . We define a distance function $D(x, y)$ which describes the amount of time necessary to travel from point $x \in M$ to $y \in M$. Then, we can describe our problem as one of traversing a minimum time path π over points in M such that we can estimate the value of $\Psi(x)$, $\forall x \in M$ to within some accuracy, ϵ .

3 Related Work

Several intelligent sampling strategies have been developed for autonomous vehicles that aim to identify hotspots, reduce resource costs, optimize sampling coverage or more accurately measure environmental phenomena. The authors in [3] explored a sampling technique, using both a team of robotic boats and static sensor nodes, in which the sensing field is partitioned into sub regions either according to equal gains or equal area and each boat is assigned a specific sub area for sampling. The readings are gathered from the static nodes and paths are computed for the boat such that they reduce the integral mean squared error. In [4], the authors explore an approach based on model parameter estimation of a variable in which the physical parameter being measured is assumed to be linearly distributed across the field and the algorithm aims to minimize the measured uncertainty in the field distribution. The algorithm also has multiple secondary objectives such as to minimize the network utility of multiple AUVs by controlling the sampling location and sampling rate using a potential function that encapsulates the network model and minimizing the energy consumption by varying the speed of the vehicle according to the energy available.

Thermoclines are believed to be an important breeding zone for marine microorganisms and hence a considerable amount of work has been done on thermocline detection and monitoring using sensor networks, gliders and other AUVs

[5, 6]. Zhang in [7] used a wireless sensor actuator network and a robot mule to detect thermoclines using distributed binary search. In this algorithm the nodes were assigned regions to sample and could move vertically by altering their buoyancy. Each node first localized the temperature variation in its own region, then combined this data with that of children nodes, forwarded it to the parent node and so on, until the final bulk data was transferred to the user. They further improved the performance of the algorithm by using a mobile robot to collect data from an active node and communicate it to another.

Sampling of Phytoplankton has also gained popularity in recent years as it plays a very important role in ocean ecology. In [8] the authors used a Dorado AUV to describe a method to detect and collect water samples at peak chlorophyll fluorescence, taking into consideration the delay in measurement while detecting the peak. The AUV followed a yo-yo pattern and used gradient following to detect a peak in the ascend stage and successfully collected the peak chlorophyll fluorescence sample at the same depth in the descend cycle.

4 Watercraft Platform

Fig. 1: Cooperative Robotic Watercraft platform.

The algorithms discussed in this paper were designed and implemented on the Cooperative Robotic Watercraft (CRW) platform [9]. CRW is a multi-robot autonomous surface vehicle, equipped with an Android smartphone that provides the inertial sensors and computing platform for the system. The CRW's design is similar to that of an airboat with a modified steering mechanism in which the entire propulsion assembly is actuated using servo motors, allowing for improved thrust vectoring, which enables sharper turns. The drive system and other electronics are interfaced to an Arduino microcontroller that communicates with the smartphone via Bluetooth. Most of the autonomy software resides on the phone while some of the application specific intelligence, such as the sampling algorithms, are implemented on a centralized operator interface that interacts with the individual vehicles via 3G or WiFi.

Water quality sensors such as dissolved oxygen, temperature, specific conductivity and pH, are mounted on the vehicle and interfaced to the system

through the Arduino. The camera on the smartphone provides real-time situational awareness about the operating environment using a steady stream of images that are processed through an image queue and displayed on the operator interface. A water sampling mechanism on board each vehicle has the ability to collect physical samples on demand for more detailed analysis in the laboratory. A diagram depicting the operation of the entire system is shown in Figure 1.

5 Bounding Filter

Rate dependent hysteresis can be observed in several sensors and poses a very serious problem to autonomous sampling where a robot continuously collects measurements while moving through a large field to create models of a physical parameter. This phenomenon cannot be ignored in dynamic sensing applications as the lag between the input and output causes a delay in the responsiveness of the sensor: for a change in input, the output of the sensor slowly and consistently approaches the actual value. This effect is not as important in static sensing applications, where the rate of change of the physical parameter in the field is much slower than the hysteresis in the sensor. However, in our case the watercraft traverses through the water while simultaneously taking measurements, making compensating for the hysteresis effect critical. Within the suite of sensors on the CRW platform, we observe this effect significantly in both temperature and dissolved oxygen.

We propose an intelligent sampling solution to this problem in the form of a filter that accounts for this effect. Rather than recording the sensed value at a location, we maintain an upper and lower bound on the predicted value in each area and use the direction of change in the sensor measurement to adjust the bounds. For example, if the gradient is trending downwards, the actual value must be lower than the value reported by the sensor, hence the upper bound can be adjusted. The inverse holds when the values are trending upwards, allowing the lower bound to be increased.

While the gradients tend to be consistent and reasonably noiseless, we use a median filter to remove occasional noise in sensor measurements. The median value over a window of readings is computed and linear regression is used to find a gradient across a larger window over the median filtered values. The gradient is then used to change either the upper or lower bound in the area if it has an absolute value above a constant defined as ϵ . Based on our initial experiments with water quality sensors exhibiting hysteresis we made two practical design decisions on the filter. Firstly, a zero gradient is the most useful gradient as it could be used to bring the upper and lower bound to the current value, since the sensor must be at the true value. However, in practice we found that this was misleading in view of discretization in the sensor output, as the gradient might appear to be zero even when it is not. Therefore we choose to ignore gradients of zero, though these will be considered in future work. Conversely, at times a sensor measurement can oscillate between two discrete levels causing an apparent gradient even when there is none. To avoid the impact of these erroneous readings on the filter, we require the gradient be above ϵ which filters out very small gradients due to oscillations in the sensor output. ϵ can be determined by analyzing a small data set obtained from a sample run. The pseudo code for the filtering process is shown below.

```

function PROCESS( $v$ )
   $Data \leftarrow v_{data}$ 
   $Windows \leftarrow MedianFilter(Data)$ 
   $gradient \leftarrow LinearRegression(Windows)$ 
   $cell \leftarrow CellFor(v_{position})$ 
  if  $gradient > \epsilon$  then
    if  $cell_{lower} < v_{data}$  then
       $cell_{lower} \leftarrow v_{data}$ 
    end if
  else if  $gradient < -\epsilon$  then
    if  $cell_{upper} > v_{data}$  then
       $cell_{upper} \leftarrow v_{data}$ 
    end if
  end if
end function

```

6 Planning With Bounding Filter

Planning for information collection is a intriguing problem that has been extensively studied in recent years [10]. Adopting the bounding filter opens up new challenges and opportunities for developing a planning algorithm. The sensor measurement with which the watercraft enters a cell is important as a change in the measurement will in turn affect the bounds of that cell. Closer the measurement with which the watercraft enters a cell is to the mid-point of bounds of the current cell, the more valuable the collected data is likely to be, since a change in the measurement will lead to the biggest expected change in the filter. Planning using the bounding filter is a challenge as it needs to take into account what measurement the sensor might output along the path, even though it can only be estimated based on the current value of the upper and lower bounds of a cell.

Upper = 50 Value = 0 Lower = 30	Upper = 60 Value = 5 Lower = 40	Upper = 70 Value = 15 Lower = 30
Upper = 60 Value = 10 Lower = 50	Sensor = 60	Upper = 90 Value = 15 Lower = 50
Upper = 70 Value = 20 Lower = 50	Upper = 90 Value = 70 Lower = 20	Upper = 100 Value = 60 Lower = 40

Fig. 2: Value of moving in different directions, given current sensor value and upper and lower bounds in adjacent areas. The bounded planner uses these values to create paths that maximize information gain.

The approach we take is shown in Figure 2 where a tree is used to expand to the most promising nodes. As paths in the tree are expanded, an estimate of the sensor value is maintained as the mid-point between the upper and lower bounds of the filter, which is then used as the expected value of the sensor as the watercraft leaves the cell. The value of going into a cell is estimated as a

function of how far the expected sensor measurement with which the watercraft enters the cell is from the mid-point between the upper and lower bounds of the cell, current difference between the upper and lower bounds and the number of times that cell has been previously visited on the path. The value of expanding a particular node in the search tree is a heuristic based on the current expansion depth and the rate that value has accumulated on the path so far. The aim of the heuristic is to encourage exploration of the paths that have the highest value.

```

function BOUNDED_PLANNER
   $n.loc \leftarrow currentlocation$ 
   $n.value \leftarrow 0$ 
   $n.sensor \leftarrow currentsensorvalue$ 
   $queue.add(n)$ 
  while  $expansions < maxExpansions$  do
     $n \leftarrow queue.poll$ 
    for  $e \leftarrow getExpansions(n)$  do
       $e.value \leftarrow n.value + max(0.0, (e.cell.upper - e.cell.lower) - ((e.cell.upper + e.cell.lower)/2) - n.sensor)$ 
       $expectedSensorChange \leftarrow smaller((e.cell.upper + e.cell.lower)/2 - n.sensor, maxChange)$ 
       $e.sensor \leftarrow n.sensor + expectedSensorChange$ 
       $queue.add(e)$ 
    end for
  end while
end function

```

We evaluated the potential of our algorithm in a simulated environment partitioned into a ten by ten grid and created a model of the dissolved oxygen sensor with mild hysteresis. The grid size is chosen based on the number of boats being used in the experiment and the size of the sensing field. The more number of grids, the better is the resolution but the longer it takes, therefore there is a tradeoff between the sensing time and resolution of sensing. Two experiments were performed to determine the utility of the filter and the planning approach. In the first experiment, the value of each cell in the environment was drawn from a uniform random distribution and in the second experiment, the value of each cell in the environment was drawn from a mixture of Gaussians. The results obtained using a simple averaging filter and the bounded filter along with three different path planning algorithms: random walk, lawnmower pattern and the bounded planner, are shown in Figures 3(a) and 3(b). In the bounded filter, we measure the error from the mid-point of the filter which may not always be a good measure of the information in the filter. For example, one of the bounds may get changed much earlier than the other due to gradients in the environment. The graphs represent the average result of 100 randomly generated environments and one simulated boat. In both cases, random movement with the bounded filter eventually leads to the lowest error. The difference is dramatically higher in the random environment as the random variation between the cells makes the hysteresis effect more crucial and therefore makes the averaging approach perform more poorly. The lawnmower pattern performs well initially, in part because measurements are collected uniformly all over the field, in turn significantly reducing the overall error. However, with the bounded filter, the lawnmower approach asymptotes towards some non-zero error as the watercraft enters cells from the same direction with the same hysteresis trend each time. The bounded planner surprisingly does not perform better than random walk

(a) Field with uniform random distribution (b) Field with mixed gaussian distribution

Fig. 3: Total error comparison in a simulated environment with different planning and filtering algorithms.

over the long run, although it has an advantage for a short period. In early stages, the good performance appears due to visiting a wider number of cells and the reason for the poorer performance later on is not completely clear and will be investigated in future work.

7 Other Algorithms

7.1 Adaptation of Level Set Approach

In some applications, a complete spatial map depicting the variation of a physical variable in the field is not required, instead the interest might be to identify an area where the concentration of that variable is above or below a certain value. For example, a fish farmer is only interested in whether any part of his pond has a dissolved oxygen level near or below what is required for the fish to survive. The level set approach can also help us identify hotspots above a certain threshold value. When only a certain level is of interest, we can adapt techniques that look specifically to find this level and isolate that area. Specifically, we have adapted an active learning approach developed by Bryan and Schneider[11]. In this approach, the observations collected to date are used to create a map of the environment and a polynomial interpolation scheme is used to estimate measurements at all locations, regardless of whether an observation has already been taken there. The marching squares algorithm is then used to generate a contour at the value of interest. Some locations along the contour may be purely based on estimates from the interpolation and some will be places where the observed values are not exactly the contour value, but the simplest contour that runs through that location.

Based on the interpolated values and estimated contour, we follow the approach by valuing points by looking at a function of uncertainty at a location and its distance to the value of interest (or contour). Given a contour value, C_v and a location value P_v , we value points as $e^{-|C_v - P_v|}$. Waypoints are assigned to locations with the highest valuation using this formula. When there is a fleet of watercraft in the same environment, a centralized planner ensures that multiple watercraft are not sent to the same location.

7.2 Adaptation of Maximum Uncertainty

As shown in our previous work [9], a simple approach to intelligent sampling is to focus on areas with maximum uncertainty. Areas that have fewer observations

Fig. 4: CRW adaptive sampling interface showing the level set and boat markers.

or where the observations have a lot of noise are locations that are significant for collecting additional measurements. Our heuristic for collecting more observations in an area is based on the variance of observations measured in that area and the number of observations taken. Specifically, the value of an observation in an area is $v_{i,j} = \sigma^2 \times \alpha^c$, where σ is the variance in the cell, $\alpha < 1.0$ and c is the number of observations taken in that cell. Values of both the variables are used since some cells may have high variance across the cell and more observations may not be capable of bringing that uncertainty down. When the boats have completed a task, the next waypoint is assigned to the nearest location which has the highest value of this function and in cases where there are multiple boats in the same environment, the entire field is sub divided into regions based on the number of boats present and each boat is assigned a specific sub region for sampling.

8 Results

In order to evaluate the algorithms, we performed analysis both in simulation as well as by extensive field trials in a variety of operational scenarios and weather conditions. Our initial experiments on dissolved oxygen monitoring were performed at Shelby fish farm and the experimental validation of algorithms were performed at Panther Hollow lake using a fleet of four CRW. The CRW was equipped with dissolved oxygen sensors from Atlas-Scientific, along with specific conductivity and temperature sensors from Decagon Devices. In order to quantify the performance of the bounded filter and planner, we calculated the total error in estimation over a period of 30 minutes and compared the results with lawn-mower pattern and maximum uncertainty sampling algorithms. For these experiments, the field was subdivided in to a set of ten by 10 grid cells and the total estimation error was computed as the sum of difference between the values of cells obtained from the algorithms and values of cells from the ground truth where we consider the ground truth as the data collected from the lake from dawn to dusk for approximately about 14 hours. We assume that the field does not vary significantly within the 30 minutes of sampling. This is a reasonable assumption for temperature and dissolved oxygen as long as there is no drastic change in the climate such as rainfall or high winds which may cause a significant change in the distribution of the physical variable in the field.

The graph in Figure 5 shows the total error comparison from sampling water temperature across Panther Hollow lake with lawn-mower pattern, highest

Fig.5: Total error comparison while sampling temperature using lawnmower pattern, highest uncertainty and bounded algorithms.

(a) Comparison between lawnmower, high-est uncertainty and bounded algorithms. (b) Comparison between measured value and filtered value from the bounded filter.

Fig. 6: Total error comparison while sampling dissolved oxygen using one boat.

uncertainty and bounded filter sampling approaches using a fleet of four CRW as well as with one CRW. For the results obtained with one CRW it can be seen that the bounded filter approach has the lowest initial error followed by gradual steps caused by the bounds being tightened significantly after the end of each plan as the planner computes the next path such that the change in sensor value is not high enough for hysteresis to be significant. From 5(b) it can be seen that even with a fleet of CRW, the bounded filter approach has the fastest convergence to the final value, proving to be scalable to various team sizes. The initial high error in the bounded planner in 5(b) is due to the selection of wider bounds in this experiment. Even though the highest uncertainty sampling has a large initial error in 5(a), the explorative nature of this algorithm forces the vehicle to sample new cells much faster and collect large amounts of samples in all the cells, therefore having the fastest convergence rate, but since the vehicle explores new cells one after the other, the hysteresis in the sensor affects the measurements in this approach considerably more, hence causing a prolonged error after the steep downward slope that can be seen in the both the plots 5(a) and 5(b). In the lawn-mower pattern sampling, as the vehicle traverses through the field in a uniform pattern, it takes a substantial amount of time for the error to reach the final convergence state as there is always the same amount of hysteresis in those specific order of cells that the vehicle follows, hence the error trend is in the form of long gradual steps.

Results from similar experiments using the dissolved oxygen sensor is shown in Figure 6(a). It can be seen that the bounded filter approach still has the lowest error compared to lawn-mower pattern and highest uncertainty sampling.

Fig. 8: Variation in dissolved oxygen observed in Panther Hollow lake.

To further investigate the error trend of the bounded filter, a plot showing the comparison between the total error in field estimation using the bounded filter and the direct sensor output is shown in Figure 6(b). The bounded filter has a lower error throughout the sampling period and converges faster than the unfiltered output, validating the novelty of this approach.

Figures 7(a) and 7(b) show the error in contour estimation when using a single boat and four boats respectively to sample temperature in the same body of water for 30 minutes. The final measured temperature distribution is thresholded at the median of the distribution to produce a binary contour ground truth estimate. This process is repeated at each timestep to produce an intermediate contour estimate, after which a logical exclusive OR operation is applied with the ground truth integrated over the sampling area to calculate the error of the contour estimate at that timestep. As evidenced in the figures, a single boat has trouble producing a converging temperature contour estimate in 30 minutes whereas four boats can easily do so for the same sampling region. This illustrates one of the main advantages of the CRW system; the multi-agent nature of the system allows for models of environmental parameters, such as temperature, to be built for large areas more quickly and reliably.

9 Conclusion and Future Work

In this paper, we address the problem of using a fleet of ASV's to accurately map different water quality parameters using a variety of intelligent sampling algorithms. We describe an adaptive autonomous sampling approach to compensate for the hysteresis that is observed in some water quality sensors using

a bounded filter and a complementary planner. The performance of this algorithm was analyzed using results from both simulation and experiments with a fleet of four CRW equipped with temperature and dissolved oxygen sensors. Our results show that the bounded filter has the least total error and converges the fastest when compared to a lawn-mower pattern and highest uncertainty sampling. Furthermore, a comparison of the results obtained with varying fleet sizes attests to the scalability of our approach consistently across experiments with both the sensors. We also describe an adaptive level set approach that can be used to identify hotspots given a certain threshold value. The error in contour estimation was computed and with four boats the error approached zero by the end of 30 minutes.

Future work on the bounded filter includes developing a method for using zero gradients to bring the bounds to the current measured value and creating a new metric that can be used to assign values to cells while evaluating the bounded filter, as currently we only take the midpoint of the bounds as the cell value at each timestep. The cooperation between the boats can also be improved by assigning the sampling location to each boat based on a heuristic such as the shortest distance to the sampling area. Investigation on the optimal fleet size to sampling region ratio, given specific convergence constraints will also be the subject of future work.

References

1. J. C. Leedekerken, M. F. Fallon, J. J. Leonard, *Mapping Complex Marine Environments with Autonomous Surface Craft*, 12th International Symposium on Experimental Robotics 2010, December 2010, New Delhi & Agra, India.
2. Hitz, G, Pomerleau, F, Garneau, M.-E and Pradalier, C et al., *Autonomous Inland Water Monitoring: Design and Application of a Surface Vessel*, IEEE Robotics & Automation Magazine, Vol. 19, Issue. 1, March 2012, pp. 62–72.
3. Zhang, Bin and Sukhatme, Gaurav S., *Adaptive Sampling for Field Reconstruction With Multiple Mobile Robots*, The Path to Autonomous Robots, Springer, 2009, pp. 1–3.
4. Dan O. Popa and Arthur C. Sanderson and Rick Komerska and Richard Blidberg and et al., *Adaptive Sampling Algorithms for Multiple Autonomous Underwater Vehicles*, IEEE/OES Autonomous Underwater Vehicles, June, 2004.
5. Cruz, N.A. and Matos, A.C., *Adaptive sampling of thermoclines with Autonomous Underwater Vehicles*, Oceans 2010, September, 2010, pp. 1–6.
6. S. Petillo, A. Balasuriya, and H. Schmidt, *Autonomous adaptive environmental assessment and feature tracking via autonomous underwater vehicles*, in Proc. IEEE Int. Conf. Oceans 2010, Sydney, Australia, May, 2010.
7. Bin Zhang and Gaurav S. Sukhatme and Aristides A. G. Requicha, *Adaptive Sampling for Marine Microorganism Monitoring*, IEEE/RSJ International Conference on Intelligent Robots and Systems, 2004, pp. 1115–1122.
8. Yanwu Zhang, Robert S. McEwen, John P. Ryan and James G. Bellingham, *An Adaptive Triggering Method for Capturing Peak Samples in a Thin Phytoplankton Layer by an Autonomous Underwater Vehicle*, Oceans, October, 2009, pp. 1–5.
9. A. Valada, P. Velagapudi, B. Kannan, C. Tomaszewski, G. Kantor and P. Scerri, *Development of a Low Cost Multi-Robot Autonomous Marine Surface Platform*, The 8th International Conference on Field and Service Robotics, Japan, July, 2012.
10. Kian Hsiang Low, John M. Dolan, and Pradeep Khosla, *Active Markov Information-Theoretic Path Planning for Robotic Environmental Sensing*, AAMAS-11, May, 2011, pp. 753–760.
11. Brent Bryan, Jeff Schneider, *Actively Learning Level-Sets of Composite Functions*, Proceedings of the 25th International Conference on Machine Learning, July, 2008.