

Persistent Ocean Monitoring with Underwater Gliders: Adapting Sampling Resolution

Ryan N. Smith*

School of Engineering Systems
Queensland University of Technology
Brisbane, QLD 4000, Australia
ryan.smith@qut.edu.au

Mac Schwager

GRASP Lab
University of Pennsylvania
Philadelphia, PA 19104, USA
schwager@seas.upenn.edu

Stephen L. Smith

Electrical and Computer Engineering
University of Waterloo
Waterloo, ON N2L 3G1, Canada
stephen.smith@uwaterloo.ca

Burton H. Jones

usCLab - The Burt Jones Group
University of Southern California
Los Angeles, CA 90089, USA
bjones@usc.edu

Daniela Rus

Distributed Robotics Lab
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
rus@csail.mit.edu

Gaurav S. Sukhatme

Robotic Embedded Systems Laboratory
University of Southern California
Los Angeles, CA 90089, USA
gaurav@usc.edu

Abstract

Ocean processes are dynamic, complex, and occur on multiple spatial and temporal scales. To obtain a synoptic view of such processes, ocean scientists collect data over long time periods. Historically, measurements were continually provided by fixed sensors, e.g., moorings, or gathered from ships. Recently, an increase in the utilization of autonomous underwater vehicles has enabled a more dynamic data acquisition approach. However, we still do not utilize the full capabilities of these vehicles. Here we present algorithms that produce persistent monitoring missions for underwater vehicles by balancing path following accuracy and sampling resolution for a given region of interest, which addresses a pressing need among ocean scientists to efficiently and effectively collect high-value data.

More specifically, this paper proposes a path planning algorithm and a speed control algorithm for underwater gliders, which together give informative trajectories for the glider to persistently monitor a patch of ocean. We optimize a cost function that blends two competing factors: maximize the information value along the path, while minimizing deviation from the planned path due to ocean currents. Speed is controlled along the planned path by adjusting the pitch angle of the underwater glider, so that higher resolution samples are collected in areas of higher information value. The resulting paths are closed circuits that can be repeatedly traversed to collect long-term ocean data in dynamic environments. The algorithms were tested during sea trials on an underwater glider operating off the coast of southern California, as well as in Monterey Bay, California. The experimental results show improvements in both data resolution and path reliability compared to previously executed sampling paths used in the respective regions.

*Corresponding Author

1 Introduction

Path planning for Autonomous Underwater Vehicles (AUVs) is required for a wide variety of applications, such as mine countermeasures, ecosystem monitoring, locating hydrothermal vents, and tracking dynamic features. For each type of deployment, a mission is generally planned *a priori* that guides the vehicle to locations of interest, or to collect a particular type of data. There are many existing methods to optimize or intelligently generate these paths, e.g., complete spatial coverage, adaptive sampling, submodular optimization, and A^* , among others. Based on the specific application, one or a combination of these techniques are used.

We propose algorithms to address the common problem in ocean science of designing a sampling¹ method to acquire data at multiple spatiotemporal resolutions to analyze ocean processes. The goal is to gather data over a predefined study area and over a long time period, where the resolution of the data varies over the space to get higher resolution in areas that are predetermined to be more interesting. Viewed over the appropriate temporal scales, such a time-series allows one to determine periodic oscillations of physical phenomena and the ocean response. We consider a given mission domain and allow a user to specify regions of high interest. A closed path for continued traversal is computed to maximize the sampling resolution in these user-defined areas, while balancing path repeatability via consideration of ocean currents. Along this path, we optimize the pitch angle of the vehicle, which changes the sampling resolution, throughout a region of interest. Repeated traversal of a regular path allows for easier assimilation of collected data with existing data or measurements, and known obstacles (e.g., shipping lanes, sea mounts, etc.) can be avoided and planned for *a priori*.

More specifically, there are four main contributions of this paper:

1. We present an algorithm to plan a closed path which passes through areas of high sensory interest, while avoiding areas that have large magnitude or highly variable ocean currents.
2. Given a path, we present an algorithm to set the pitch angle at which the glider moves along the path to ensure higher sample density is achieved in areas of higher scientific interest.
3. The two algorithms are then combined into a single, iterative algorithm that is shown to converge to a lower cost plan than either of the algorithms acting alone.
4. We validate our algorithms by implementing the computed strategies on an autonomous underwater glider. We present three implementations of our technique which balance persistent monitoring to resolve large-scale events, with the collection of high-resolution data along the designed path. The specific ocean phenomenon that we study using these algorithms is the occurrence and life-cycle of harmful algal blooms.

Field trials were carried out in both the Southern California Bight (SCB)² and in Monterey Bay, California (MB)³, with a motivation to understand the connections between small-scale biophysical processes and large-scale events common to both regions. Algal blooms are the large-scale events of interest, and especially those blooms composed of toxin producing species, commonly referred to as Harmful Algal Blooms (HABs) (Anderson et al., 2000; Ramsdell et al., 2005; Pitcher et al., 2005; Hoagland and Scatista, 2006). On a smaller scale, we are interested in the environmental triggers leading to the onset, evolution and ultimate mortality of HAB events. These triggers are poorly understood, but are commonly attributed to changes in physical structure and dynamics, and the nutrient flux from upwelling (Ryan et al., 2009); pressure gradients leading to the propagation of internal waves (Noble et al., 2009); the vertical fluxes of phytoplankton and

¹In this paper, we use sampling in a general context meaning the gathering or collection of sensor data; not the acquisition of water that is commonly implied in ocean science literature.

²The SCB is the oceanic region contained within 32° N to 34.5° N and -117° E to -121° E.

³Monterey Bay is the oceanic region contained within 36.5° N to 37° N and -121.75° E to -122.25° E.

toxins through the water column (Sekula-Wood et al., 2009); and anthropogenic inputs (DiGiacomo et al., 2004; Corcoran et al., 2010; Cetinic et al., 2010).

Various processes affect our approach to glider deployment patterns. In coastal regions such as described in this paper, because of the coastal boundary, the primary current directions are along coast both in the upper and lower layers (Hamilton et al., 2006). This effectively lengthens the alongshore scales of variability. In contrast cross-shelf scales of variability may be much greater and scaled by the shelf width and topography. Processes such as internal tides and internal waves are often formed by the topography and tend to propagate cross-shelf resulting in significant cross-shelf variability. In addition, coastal inputs such as river discharges are constrained in the offshore direction by the coastal currents, have extensive alongshelf scales, and shallow vertical scales (Warrick et al., 2007; Washburn et al., 2003). On the large scale, coastal California is characterized by baroclinic Rossby radii of deformation that range from 30 km in the south to ~ 20 km in the north (Chelton et al., 1998). In a region like southern California where the currents are more dependent on the alongshore pressure gradient than the local wind forcing regime and where the topography is more important to constraining the flow, this is probably less important than on an open coastline. Various scales of features have been observed in this region. Coastal plumes from freshwater runoff typically have cross-shelf scales on the order of kilometers (Washburn et al., 2003), and submerged outfall plumes may have cross-shelf scales that are 1 – 2 km (Jones et al., 2002). Phytoplankton patches, both surface and subsurface, may be more spatially extensive in the large scale (Thomas et al., 2009), but have small scale variability in the coastal region (Schnetzer et al., 2007).

1.1 Previous Work

There is a significant amount of literature on planning information-rich, and adaptive paths for environmental monitoring, however existing techniques do not address the problem posed above, while additionally incorporating the unique constraints and capabilities of autonomous gliders. Here, we want to take advantage of a glider’s ability to perform the persistent, long-term surveillance necessary for studying ocean processes, and provide missions that consider their low navigational accuracy and a limited means of control.

In the literature, one method related to our problem is to plan covering paths over the environment. Examples of this appear in agriculture (Oksanen, 2007), (Taïx et al., 2006), general robotics (Choset, 2000), and AUV applications (Anstee and Fang, 2010). These methods rely on a nontrivial finite sensor footprint, whereas the typical suite of in situ sensors on-board a glider take only point measurements, implying that a complete, synoptical, spatial coverage approach is infeasible. As previously mentioned, horizontal variability in this region is not as significant as vertical variability, i.e., it is more important to obtain complete coverage during a vertical profile than it is over a horizontal transect, which implies that complete horizontal coverage is not necessary. Other approaches to surveillance and monitoring with AUVs essentially boil down to achieving the best estimation or re-creation of a scalar field via intelligent planning or adaptive sampling (Singh et al., 2009; Paley et al., 2008; McGann et al., 2008; Leonard et al., 2010). These methods can present a biased optimization problem. Specifically, an AUV operates multiple sensors that simultaneously gather measurements along the trajectory of the vehicle. If we use a notion of path planning with the intent to reconstruct or estimate a scalar field, the question arises as to which data define the scalar field, i.e., do we consider temperature, salinity or chlorophyll as the dominant measurement? In practice, the relevant parameters for a specific mission are science-driven, and clearly defined by the process under investigation. To computing a general, long-term, persistent monitoring mission, collecting data relating to multiple ocean processes, one may naively choose to optimize over all scalar fields of collected data via some weighted cost function. However, correctly defining these weights, and the relationship between the collected data is a difficult problem and an open question, and generally the primary reason for the impetus to collect data within the region. Also, data may be uncorrelated both in space and time, thus defining the *best* sampling path based on static measurements could lead to an ill-posed optimization. Finally, there has recently been work on persistent monitoring (Smith and Rus, 2010), (Smith et al., 2010e), where the frequency of visits to each region is adapted to the time scale on which that region changes. This study draws from the two later references, and is directed towards developing a sampling strategy that is sensitive to the relative importance




Figure 1: A schematic example of a computed glider path for a region Q . The regions of interest Q_i are colored based on their user-defined importance p_i . The intersection of region Q_i with segment γ_j is given by l_{ij} . The proposed optimization algorithm in Section 5 produces a set of pitch angles ϕ_j for the glider to implement along each segment γ_j to achieve an appropriate sampling resolution.

of different regions, while being naturally amenable to the unique constraints of the autonomous underwater glider.

We organize the remainder of this paper in the following way. The next section provides a formal problem description and describes the autonomous underwater glider that was deployed for the experiments. The particular constraints of the underwater glider are necessary to motivate the path planning and pitch optimization algorithms. Section 3 explains the regional ocean models used for considering ocean currents. In Section 4, we present the Zig-Zag in the Tranquill Ocean Path Planner (ZZTOPP) algorithm (Smith et al., 2011), a path planner for designing a repeatable path for long-term AUV data collection and ocean monitoring. Section 5 defines an algorithm that varies the sampling resolution along the path generated from the algorithm presented in Section 4. We follow this by combining these two algorithms into a single, iterative algorithm in Section 5.1. In Section 6 we motivate our field trials through a presentation of scientific details regarding the ocean science application addressed by this study. This leads into the presentation of results and analysis from the implementation of our computed missions during sea trials in Section 7. A summary of results and observations, along with areas of future investigation and extensions are included in Section 8.

2 Problem Set-up and Vehicle Description

We consider a mission domain $Q \subset \mathbb{R}^2$. Within Q there are $n - 1$ user-defined regions of interest $Q_i \subset Q$, where $i \in \{1, \dots, n - 1\}$. We define an n^{th} region $Q_n = Q \setminus \cup_{i=1}^{n-1} Q_i$ to be the *background*. Thus, we have $Q = \cup_{i=1}^n Q_i$. Each region Q_i , $i \in \{1, \dots, n\}$ is also assigned an importance level, $p_i \in [0, 1]$, by the user. We assume that the regions Q_1, \dots, Q_{n-1} and their weights, p_1, \dots, p_n are specified by an ocean scientist, and the weights p_i determine the sampling importance of a region relative to all other regions. This notation is presented graphically in Figure 1.

Problem Statement. *Design a sampling path γ to be continually executed by an autonomous glider, which provides a synoptic view of Q , with the ability to gather data at multiple spatial resolutions based the relative importance of each $Q_i \subset Q$.*

Our proposed method for this problem is to compute a closed-path γ consisting of $0 < m < \infty$ waypoints that steers the glider through the designated regions $Q_i, i \in \{1, \dots, n-1\}$ while avoiding areas of strong or highly variable ocean currents, which may decrease path repeatability and accuracy. Along the computed path, we optimize the pitch angle of the vehicle to alter the sampling resolution.

For the mission posed in the Problem Statement, a first choice may be a lawnmower-type path. This back-and-forth pattern is a widely-used path plan for many applications in robotics. Optimizing this type of path for a given application involves varying the length of the legs to satisfy the constraints or sampling requirements of the mission. For vehicles with adequate localization and navigational accuracy a lawnmower path can be executed quite well. For implementations onto autonomous gliders, an *optimized* lawnmower-type path may not be the best choice as a design for a repeatable mission. Since gliders are dead-reckoning vehicles, their navigational accuracy is understandably poor; in the absence of currents we expect a three-sigma, cross-track error of ~ 600 m for a 2 km trajectory, see (Smith et al., 2010c). Assuming that the glider is not surfacing every hour, and we are not in an area void of ocean currents, traversing the long legs of the lawnmower path can result in large navigational errors. Depending on the size of the region of interest and the chosen plan, these errors could be larger than the shorter, *turn-around* segments of the lawnmower path. This will result in either, the glider surfacing twice in a short time period (wasting time on the surface), or the glider trivially achieving both waypoints with a single surfacing (not executing the short leg of the lawnmower path at all). Considering an optimal lawnmower path, one must be careful that the shorter legs do not become too short, else the poor navigation of a glider will not execute them properly and the planning effort was wasted. For example, Figure 2 displays a planned lawnmower path in black, and the actual executed path by a glider in white. This was an actual deployment off the northwest tip of Catalina Island, CA. This region experiences heavy ship traffic, and as a result, surfacings are limited to 8 hour intervals for safety reasons. Additionally, as was the case during this experiment, eddies spin up in this region causing a very complex and unpredictable current regime. Eddy-type currents render the glider’s on-board current correction algorithm minimally helpful. Since the correction is based on depth-averaged currents from the previously executed segment, the algorithm will not predict a circular current well, as can be seen in Figure 2. The experimental result presented in Figure 2 is close to a worst case scenario, but illustrates the type of implementation issues that can occur without advance planning, and consideration for the region of the deployment. One item to note from this experiment is that the executed path of the glider resembles a zig-zag structure that, under the appropriate circumstances, may provide the same synoptic spatial coverage of the area as the planned lawnmower path. Additionally, a zig-zag-type path is more forgiving with respect to implementation by a glider. Given this information, and that the reference path presented in Figure 9 has a zig-zag structure, as opposed to a lawnmower pattern, we chose to consider a zig-zag-shaped path that broadly covers a given region of interest for our base path structure.

2.1 Autonomous Underwater Vehicle

The vehicle used in this study is a Webb Slocum autonomous underwater glider (Webb Research Corporation, 2008), see Figure 2.1. A Slocum glider is a 1.5 m (length) by 21.3 cm (diameter), torpedo-shaped vehicle designed for long-term (~ 1 month) ocean sampling and monitoring (Schofield et al., 2007), (Griffiths et al., 2007). Originally conceived by Henry Stommel (Stommel, 1989), the autonomous glider has become an integral tool of ocean science in this decade.

Underwater gliders such as the Webb Slocum glider traverse piecewise linear paths defined by waypoints. At each waypoint they surface, orient themselves in the direction of the next way point using GPS, and submerge. While moving to the next waypoint, they glide up and down in a sawtooth pattern of fixed pitch, taking samples on the descending edge of each sawtooth. The spatial resolution of the data collected along a leg of the path is determined by the pitch of the sawtooth that the glider traverses— the steeper the pitch, the higher the data resolution. Furthermore, the glider operates in a dead-reckoning fashion along each leg of the path, and is constrained to use a constant sawtooth pitch along that leg.

Underwater gliders *fly* through the water driven entirely by a variable buoyancy system rather than active




Figure 2: Planned lawnmower pattern (black waypoints and path) and the actual executed path (white surfacing locations and path) for a glider deployment off of the northern tip of Santa Catalina Island, CA.

propulsion. Wings convert the buoyancy-dependent vertical motion into forward velocity. Due to this method of locomotion, gliders are typically operated as slow moving AUVs with operational velocities on the same order of magnitude as oceanic currents (~ 30 cm/s). By changing the amount of mass exchanged between the inside and outside of the hull, one can alter a glider's velocity. In particular, larger gliders have attained across-the-water speeds rivalling that of some thruster powered AUVs. For specific operational details of Slocum gliders, see e.g., (Webb Research Corporation, 2008) or (Schofield et al., 2007). In addition, gliders are not particularly accurate in the execution of a prescribed path. Considerable work has been done on low-level control of underwater gliders, see (Leonard and Graver, 2001; Graver, 2005), which we do not address here. The authors have studied methods to increase the navigational accuracy of autonomous gliders by use of ocean model predictions, see (Smith et al., 2010c; Smith et al., 2010d). This work provides a starting point, and motivates study in planning methods that reduce uncertainty of the forces experienced from ocean currents.

Between pre-programmed surfacings for GPS fixes and data transfer, the glider dead reckons its position using a magnetic compass, depth sensor, and altimeter. Briefly, an example mission for a standard glider consists of a set maximum depth along with an ordered list of geographical waypoints (W_1, \dots, W_n). An exact path or trajectory connecting these locations is not prescribed by the operator, nor are the controls to realize the final destination. When navigating to a new waypoint, the present location W of the vehicle is compared to the next prescribed waypoint in the mission file (W_i), and the on-board computer computes a bearing and range for execution of the next segment of the mission. We will refer to the geographical location at the extent of the computed bearing and range from W to be the aiming point A_i . The vehicle then dead reckons with the computed bearing and range towards A_i with the intent of surfacing at W_i . The glider operates under closed-loop heading and pitch control only. Thus, the computed bearing is not altered, and the glider must surface to make any corrections or modifications to its trajectory. When the glider completes the computed segment (i.e., determines that it has traveled the requested range at the



Figure 3: One of two Slocum gliders owned and operated by the USC Center for Integrated Networked Aquatic PlatformS (CINAPS) (Smith et al., 2010b) preparing to start a mission off the Northeast coast of Santa Catalina Island, CA.

specified bearing), it surfaces and acquires a GPS fix. Regardless of where the vehicle surfaces, waypoint W_i is determined to be *achieved*. The geographic positional error between the actual surfacing location and W_i is computed, and any error between these two is fully attributed to environmental disturbances (i.e., ocean currents). A depth-averaged current vector is computed, and this is considered when computing the range and bearing to W_{i+1} , the next waypoint in the mission list. Hence, A_i is in general not in the same physical location as W_i . The offset between A_i and W_i is determined by the average velocity and the perceived current experienced during the previous segment. For the experiments presented in the sequel, this on-board, current-compensation algorithm was utilized by the glider. Work is ongoing to integrate this algorithm with the predictive, current-compensation algorithm presented in (Smith et al., 2010a).

3 Navigating in an Uncertain Environment

Dealing with repeatability and minimizing uncertainty along a path are complex issues when considering an underwater vehicle. The ocean is a highly dynamic, time-varying, nonlinear system that imparts large magnitude external forces upon a vehicle. Some of these forces and moments, like buoyancy and viscous damping, can be estimated with reasonable accuracy for a given vehicle, e.g., see (Allmendinger, 1990; Fossen, 1994). However, the forces and moments related to ocean currents can greatly affect the navigational accuracy of a dead reckoning glider (Smith et al., 2010a; Smith et al., 2010c). Following the ideas presented in these references, we utilize ocean model outputs to assist in planning and navigation.

In this study, we consider two approaches to incorporate ocean model outputs into a path planning algorithm. An initial, conservative approach, as presented in (Smith et al., 2011), we consider regions with the largest magnitude predicted currents, averaged over a given time frame. A path planner is rewarded for avoiding these strong current areas. This is conservative, since we steer the vehicle away from areas that experience larger currents on average. Details of the computations and implementation of this strategy are presented in Section 7.2. An alternative approach is to not simply consider the magnitude of the current, but to instead examine the temporal variability of the current. This is a more aggressive approach since we might steer the glider through areas of strong currents. However, if these strong currents have low variability over time, we can predict them with better accuracy, and can apply the algorithms presented in (Smith et al., 2010a)

to achieve good navigational accuracy. Determining the variability of a scalar field can be done in many ways. For our application, we use the variance of the current speed as predicted by standard regional ocean models.

For this study, two different predictive ocean models were used; one for operations in the SCB and one for operations in MB. For the SCB, we utilize data from the Regional Ocean Model System managed by the Jet Propulsion Laboratory (JPL), California Institute of Technology. For experiments in MB, we utilized the Navy Coastal Ocean Model (NCOM) Innovative Coastal-Ocean Observing Network (ICON) managed and run by the United States Naval Research Laboratory (NRL) field site in Monterey, CA. Below is a brief overview of each of the two models. For specific details, we refer the interested reader to the cited publications and the references contained therein.

3.1 ROMS - Southern California Bight

The Regional Ocean Model System (ROMS) run at the Jet Propulsion Laboratory, California Institute of Technology is a split-explicit, free-surface, topography-following-coordinate oceanic model. The model output has three nested horizontal resolutions covering the U.S. west coastal ocean (15 km), the southern California coastal ocean (5 km) and the SCB (2.2 km). The three nested ROMS domains are coupled online and run simultaneously exchanging boundary conditions at every time step of the coarser resolution domain. ROMS provides hindcasts, nowcasts and hourly forecasts (up to 72 hours) for the SCB, (Vu, 2008), (Li et al., 2008). The operational model assimilates temperature and salinity data from autonomous vehicles, sea surface temperature from satellites, and surface currents from the high-frequency radar network. Detailed information on ROMS can be found in (Shchepetkin and McWilliams, 2005) and (Shchepetkin and McWilliams, 1998).

3.2 NCOM ICON - Monterey Bay

The Navy Coastal Ocean Model (NCOM) run by the Naval Research Laboratory is a primitive equation, 3D, hydrostatic model that uses the MellorYamada level 2.5 turbulence closure scheme, and the Smagorinsky formulation for horizontal mixing. This model is a hierarchy of different resolution models for the West Coast of the United States (Shulman et al., 2007). The high-resolution model used here is the NCOM Innovative Coastal-Ocean Observing Network (ICON) (Shulman et al., 2007), and it is set up on a curvilinear orthogonal grid with horizontal resolution of approximately 0.4 km. There are 30 sigma coordinate vertical levels. Assimilation of temperature and salinity data provided from autonomous platforms is performed every 12 hours. Model outputs contain a 24 hour hindcast, 24 hour nowcast and 24 hour forecast each day. For more details on this model, please see (Shulman et al., 2009) and (Shulman et al., 2010).

3.3 Current Speed Variance

Using historical model outputs of the speed of the current from the two ocean models, we calculated the variance of the current speed at each point in a grid over the region of interest. We begin with the four-dimensional, time-series, outputs from the ocean models, which give the northward, $V(x, y, z, t)$, and eastward, $U(x, y, z, t)$, components of the ocean current for each hour in a three-dimensional grid. Here, $t \in [0, 23]$ hours, $z \in [0, 80]$ m, and x and y are the longitude and latitude grid locations, respectively, that cover the designated area of interest. First, we average the U and V components over the depth and then compute the speed at each latitude and longitude position for each time, giving

$$S(x, y, t) = \left(\left(\frac{1}{N_z} \sum_{i=1}^{N_z} U(x, y, z_i) \right)^2 + \left(\frac{1}{N_z} \sum_{i=1}^{N_z} V(x, y, z_i, t) \right)^2 \right)^{1/2},$$




Figure 4: Ocean current variance for Monterey Bay region of interest for three weeks of data. The color bar scale has the units m^2/s^2 .

where N_z is the number of depth points in the grid.⁴ Without loss of generality, we assume $S(x, y, t)$ has zero temporal mean, i.e., any initially nonzero temporal average has already been subtracted out. One needs to be careful, however, that the temporal scale over which the data are averaged captures the proper variability for the scale of the process being examined. Finally, we commute the variance at each point (x, y) to be

$$\sigma^2(x, y) = \frac{1}{N_t - 1} \sum_{i=1}^{N_t} S(x, y)^2,$$

where N_t is the number of time samples. The sum is divided by $N_t - 1$ rather than N_t because, if one considers the speeds at each time to be independent identically distributed random variables, this gives an unbiased estimate of their variance. The current variance over three weeks of data for the MB region of interest is shown in Figure 4, and the variance from two months of data for the SCB is shown in Figure 5.

The above defined process reduces vector information into a scalar speed. While we would like the proposed method to be as general as possible, care needs to be taken in applying this to a general region of interest. The initial motivation here stems from the use of depth averaged currents by the glider’s on-board current correction algorithm, and with the understanding that the primary deployment location is within the SCB. This reduction is sensible for the SCB because tidal currents along the southern California coast are mixed tides that have an M2 component (period of 12.42 hours), with a significant fortnightly spring-neap cycle. Nearshore, the tidal signatures affect the modulation of the alongshore currents. Decorrelation times for tidally-induced current and velocity components are approximately 5 days (Bratkovich, 1985). Regions in which the tides are dominated by the M2 component may require careful examination to ensure that this reduction assumption remains valid.

Next, we present the two algorithms that together are used to design a sampling strategy for persistent and adaptive monitoring of an oceanic region with designated areas of *high-interest*. The path planner presented in Section 4 details an optimization that rewards visiting each Q_i and penalizes navigation through areas with large mean magnitude or high variability in ocean current, while respecting an overall time budget for the traversal of the path. Given this path, the algorithm presented in Section 5 optimizes the pitch angle

⁴Based on the trajectory of an autonomous glider and the resolution of current ocean models, one may assume that the vehicle spends equal time at each depth along the sawtooth and thus is affected by depth-integrated currents.




Figure 5: Ocean current variance for the SCB region of interest for two months of data. The color bar scale has the units m^2/s^2 .

along each segment to obtain an adaptive sampling resolution for each Q_i , while also maintaining the overall time budget for the traversal of the path.

4 Path Planning

Here we describe the algorithm used to plan the path of the robot, which we have named the Zig-Zag in the Tranquil Ocean Path Planner (ZZTOPP). The planner uses a constrained version of the Bellman-Ford algorithm (Cormen et al., 2001) to find the approximate maximum reward path through a graph of possible waypoints. We formulate a reward function that balances the desire to avoid high current (or high current variance) regions, with the desire to move through areas where sampling is most important. We construct the graph of possible waypoints by dividing the environment into a given number of subsections on alternating sides of a dividing line. The dividing line specifies the main axis of orientation for the path. A graph is generated consisting of edges between each of the points in successive subsections. One waypoint is then selected from each of the successive subsections to make up the path. We constrain the path to take no longer than a given time to complete, and this constraint makes the path optimization NP-hard (Jüttner et al., 2001). We therefore use a Lagrangian relaxation version of the Bellman-Ford algorithm to approximately solve the constrained optimization in a reasonable time, with known tight approximation bounds.

The above method of path construction gives us the desired zig-zag shape for the path with transects crossing the region of interest. We constrained the path to be constructed in this way for several practical reasons. Firstly, paths that are hand designed by ocean scientists have this general zig-zag shape (see for example the cyan path in Figure 9). Keeping this shape is important so that the automatically generated paths look reasonable to ocean scientists, and so that they can be fairly compared with paths that are already in use. Secondly, as we have already argued, the zig-zag shape is preferable to other common choices (e.g. a square lawnmower shape as in Figure 2) for path accuracy. Finally, it is desirable for the scientist to specify the main axis of orientation for the path to ensure that the glider traverses important known features, such as shelves or canyons, in the right orientation.

The ZZTOPP algorithm requires the following inputs: 1) the mission domain Q ; 2) the high interest regions

Q_i , with associated weights p_i , $i \in \{1, \dots, n\}$; 3) a time budget T , giving the maximum allowable time for the robot to do a complete cycle of the path; 4) the desired number of waypoints $m \in M < \infty$; and 5) a line L dividing Q into two halves, and defining the ordinate axis of the proposed zig-zag path. Here M represents the entire set of possible waypoints, which is determined by intersecting Q with the discretized grid of the considered ocean model, i.e., ROMS or NCOM. This provides a discretized set of candidate waypoints based on the resolution of the underlying model. As the models considered are on the order of magnitude of the navigational accuracy of the glider, a finer resolution discretization would not provide a significant difference in the executable paths.

Given m , let $\{w_1, \dots, w_m\}$ be the m waypoints that define the path γ , and for each $j \in \{1, \dots, m\}$, let γ_j be the line segment connecting the endpoints w_j and w_{j+1} , where we let $w_{m+1} := w_1$. Then, given L , this algorithm generates an alternating *checkerboard* pattern consisting of m cells. This is depicted in Figure 8. We restrict waypoint w_j to lie in checkerboard cell j . The path is created by choosing the waypoint in each cell of the checkerboard, such that it minimizes a cost function capturing both a penalty for high ocean currents or high variability and a reward for traversing through high-interest regions.

We evaluate the effect of ocean currents as follows. Let $M \subset Q$ denote the finite set of points that define the discretized output from the regional ocean model. We use a value $\nu(q)$ for each $q \in M$, representing the magnitude or variance of ocean currents at q . We can interpolate this data to obtain a function $\nu : Q \rightarrow \mathbb{R}_{\geq 0}$, which is defined at each point $q \in Q$. (In what follows, for each point $q \notin M$, we assign the value of the nearest neighbor in M .) Then, the cost of a path due to ocean currents is given by

$$\sum_{j=1}^m \int_{\gamma_j} \nu(q) dq. \quad (1)$$

The reward for passing through regions Q_i is defined as follows. Define the length of the intersection of a region Q_i with a path segment γ_j by $l_{ij} = |Q_i \cap \gamma_j|$, where $|\cdot|$ denotes the length of the segment. For a given path γ , the length of γ passing through region Q_i is $\sum_{j=1}^m l_{ij}$. Since we want to spend more time in regions of higher interest, we define the reward as

$$\sum_{i=1}^n p_i \sum_{j=1}^m l_{ij}, \quad (2)$$

where the p_i 's are the weights assigned by a scientist. Therefore, the cost of a set of waypoints $W = \{w_1, \dots, w_m\}$ is defined as

$$H(W, \lambda) := \lambda \sum_{j=1}^m \int_{\gamma_j} \nu(q) dq - (1 - \lambda) \sum_{i=1}^n p_i \sum_{j=1}^m l_{ij}, \quad (3)$$

where $\lambda \in [0, 1]$. For $\lambda = 1$ we only consider ocean currents, and for $\lambda = 0$ we consider only high-interest regions. A specific $\lambda \in [0, 1]$ is chosen to equivocate the order of magnitude of the mean magnitude or current variation input with the weighting of the regions of interest in the cost function. Given $\lambda \in [0, 1]$, we search all possible sets of m waypoints for the set W^* that minimizes $H(W, \lambda)$.

The final component of the optimization is the constraint T on the time to complete one full cycle of the path. Letting $v_{\text{hor},j}$ be the speed over ground of the glider along segment γ_j , the total time to complete the path is given by

$$\sum_{j=1}^m \frac{|\gamma_j|}{v_{\text{hor},j}}. \quad (4)$$

We have performed experiments to determine $v_{\text{hor},j}$ as a function of the pitch angle of the glider, ϕ_j , for the operating regime of pitch angles $15^\circ \leq \phi_j \leq 35^\circ$. By performing a least-squares fit, we obtain a relationship of

$$v_{\text{hor},j} = ax_j + b, \quad (5)$$

where $x_j = \tan \phi_j$, $a := -0.05 \text{ m/s}$, and $b := 0.275 \text{ m/s}$. The coefficient of determination (or R^2 value) is 0.986 indicating a good linear fit.

Then, for a fixed λ , our optimization is

$$\begin{aligned} & \text{minimize } H(W, \lambda) \text{ over } W \\ & \text{subject to } \sum_{j=1}^m \frac{|\gamma_j|}{ax_j + b} \leq T \end{aligned}$$

Our method for performing the optimization is to discretize Q by considering only $Q \cap M$, so that each checkerboard region contains a finite number of candidate waypoints. Then, we define a graph whose vertices are these discretized points, and whose edges connect waypoints in checkerboard square j to waypoints in square $j + 1$. This defines a directed graph, where every cycle is of the form w_1, \dots, w_m, w_1 . We define two weights on the edge γ_j connecting w_j to w_{j+1} . The first weight gives the contribution of edge γ_j to the cost $H(W, \lambda)$,

$$c_1(\gamma_j) = \lambda \int_{\gamma_j} \nu(q) dq - (1 - \lambda) \sum_{i=1}^n p_i l_{ij}. \quad (6)$$

The second weight gives the travel time

$$c_2(\gamma_j) = \frac{|\gamma_j|}{ax_j + b}. \quad (7)$$

Note that the weights c_2 are non-negative, while the weights c_1 may take on negative values.

By fixing the waypoint w_1 , the optimization problem becomes a constrained shortest path problem (with possibly negative edge weights): minimize $\sum_{j=1}^m c_1(\gamma_j)$ subject to $\sum_{j=1}^m c_2(\gamma_j) \leq T$. An exact solution to this problem is known to be NP-hard (Jüttner et al., 2001). However, there exist very good heuristics for finding approximate solutions. In this paper we utilize the heuristic based on Lagrangian relaxation of the constraint (Jüttner et al., 2001). The relaxed problem can be solved by repeatedly using the Bellman-Ford algorithm. The algorithm requires $O(|E| \log^3 |E|)$ iterations. In the worst-case, an iteration of the Bellman-Ford algorithm runs in $O(|V||E|)$ time, and thus heuristic computes a constrained solution in $O(|V||E|^2 \log^3 |E|)$ time. We solve the optimization for each w_1 , selecting the position w_1 with minimum cost. We determine an approximate minimizer $H(W, \lambda)$. The above description is summarized in Alg. 1.

Given that there are N_j points in checkerboard cell j , the total number of vertices in the graph is $N := \sum_{j=1}^m N_j$. The number of edges in the graph is $\sum_{j=1}^m N_j N_{j+1}$, where $N_{m+1} := N_1$. Since the Lagrangian relaxation algorithm runs in worst-case time of $O(|V||E|^2 \log^3 |E|)$, the ZZTOPP algorithm runs in $O(N_1 N (\sum_{j=1}^m N_j N_{j+1})^2 \log^3 (\sum_{j=1}^m N_j N_{j+1}))$ computation time. If $N_j = \bar{N}$ for each j , then this simplifies to $O(m \bar{N}^6 \log^3 \bar{N})$.

We can readily extended this algorithm to automatically search over choices for the main dividing line L . Suppose that the orientation of the path is not important to the ocean scientist. Let us specify that L passes through the centroid of the mission domain, Q , and consider a set of possible orientation angles for L between 0 and 2π . Suppose there are N_L angles in this set. Then the ZZTOP algorithm can be run at each orientation, and the one with the lowest cost path can be selected. This will, of course, increase the running time by a factor of N_L

5 Sample Resolution Optimization

Once the path is planned using the algorithm above, we use the following algorithm to determine the pitch angles along the path to ensure higher sample resolution in more important areas, and lower sample resolution in less important areas. The sampling resolution for a glider is altered by changing the pitch angle of the sawtooth pattern that the glider executes as it traverses the segments of the given path.

Algorithm 1: Zig-Zag in the Tranquil Ocean Path Planner (ZZTOPP)

Input : 1) The high interest regions Q_i and their associated importance levels p_i ; 2) the number of waypoints m ; 3) the axis L ; and 4) the parameter $\lambda \in [0, 1]$.

Output: A set W of m waypoints which seeks to minimize the cost function $H(W, \lambda)$ subject to the time constraint T .

- 1 Compute the checkerboard regions from the line L and number of waypoints m .
- 2 Determine the function $\nu : Q \rightarrow \mathbb{R}_{\geq 0}$ defining the average magnitude or variability of the ocean current from the regional ocean model data.
- 3 Discretize checkerboard regions.
- 4 Generate a graph G with vertices given by discretized points, edges connecting each point in checkerboard region j to each point in checkerboard region $j + 1$.
- 5 Compute two sets of edge weights. For edge γ_j connecting w_j and w_{j+1} , the weights are

$$c_1(\gamma_j) = \lambda \int_{\gamma_j} \nu(q) dq - (1 - \lambda) \sum_{i=1}^n p_i l_{ij}$$

$$c_2(\gamma_j) = \frac{|\gamma_j|}{ax_j + b}.$$

6 Set BestScore $\leftarrow +\infty$

7 **foreach** candidate waypoint w_1 **do**

8 Compute cycle W_{cand} in G containing w_1 that minimizes $\sum_{j=1}^m c_1(\gamma_j)$ subject to $\sum_{j=1}^m c_2(\gamma_j) \leq T$.

9 If $H(W_{\text{cand}}, \lambda) < \text{BestScore}$ then set $W \leftarrow W_{\text{cand}}$ and BestScore $\leftarrow H(W_{\text{cand}}, \lambda)$.

10 **Output** W .

The behavior of a glider on segment γ_j , $j \in \{1, \dots, m\}$ is controlled by the following quantities:

1. pitch angle $\phi_j \in [\phi_j^{\min}, \phi_j^{\max}]$
2. minimum depth d_j^{\min} , and
3. maximum depth d_j^{\max} .

Define the sample density in a region Q_i , denoted by ϱ_{Q_i} , to be the number of samples taken in region Q_i at a given depth. Similarly, we define the sample density along a segment γ_j , denoted by ϱ_{γ_j} , to be the number of samples along the segment taken at a given depth. Each time the glider descends along the edge of a tooth, it takes at most one sample at a given depth d , (to avoid hysteresis effects, the glider does not take measurements while ascending), as shown in Figure 6. The sample density ϱ_j along segment γ_j , assuming an idealized triangular sawtooth (ignoring currents, disturbances, hydrodynamics effects, etc.), is found from simple geometry to be

$$\varrho_{\gamma_j} = \frac{\tan \phi_j}{2(d_j^{\max} - d_j^{\min})}. \quad (8)$$

Then we compute the sample density in a region Q_i as the number of samples in that Region Q_1 at fixed depth,

$$\varrho_{Q_i} = \sum_{j=1}^m l_{ij} \varrho_{\gamma_j}.$$

We propose that the pitch angles ϕ_1, \dots, ϕ_m be set so as to maximize a measure of the total sampling reward. The optimization of the pitch angles is subject to constraints on the minimum and maximum pitch, as well




Figure 6: The sawtooth pattern that the glider follows along each path segment γ_j is shown with its relevant parameters labeled. The sawtooth pitch ϕ_j controls the sample density, which we define as the number of samples per distance at a fixed depth d .

as on the total time to complete one cycle of the path. To motivate our form for sampling reward, consider a region Q_i and its sampling resolution ϱ_{Q_i} . Naturally, in increasing the sampling resolution ϱ_{Q_i} , we increase the “reward.” However, it is intuitive that this reward will be subject to *diminishing returns*. That is, for a given region, the marginal reward of additional samples decreases as the total number of samples increases. This can be captured via a concave, monotonic function $G : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, satisfying $G(0) = 0$. In the experiments section $H(z) = \sqrt{z}$.

We define the total reward of a set of pitch angles as

$$C(\phi_1, \dots, \phi_m) := \sum_{i=1}^n p_i G(\varrho_{Q_i}). \quad (9)$$

Since G is concave, the reward C is also concave. To compress notation, we define

$$\beta_{ij} = \frac{l_{ij}}{2(d_j^{\max} - d_j^{\min})}, \quad \beta_i = [\beta_{i1} \ \dots \ \beta_{im}]^T, \quad (10)$$

$$x_j = \tan \phi_j, \quad \text{and} \quad x = [x_1 \ \dots \ x_m]^T, \quad (11)$$

so that $\varrho_{Q_i} = \beta_i^T x$. Also, define the vectors

$$x_j^{\min} = \tan \phi_j^{\min}, \quad x^{\min} = [x_1^{\min} \ \dots \ x_m^{\min}]^T, \\ \text{and} \quad x_j^{\max} = \tan \phi_j^{\max}, \quad x^{\max} = [x_1^{\max} \ \dots \ x_m^{\max}]^T.$$

With the above definitions, the reward in (9) is written as

$$C(x) = \sum_{i=1}^n p_i G(\beta_i^T x).$$

The final component of the optimization is the time constraint. From Section 4, the time to complete the path is given by

$$\sum_{j=1}^m \frac{|\gamma_j|}{v_{\text{hor},j}}.$$

where, as mentioned previously, we have experimentally obtained the relation

$$v_{\text{hor},j} = ax_j + b, \quad (12)$$

where $x_j = \tan \phi_j$, $a := -0.05 \text{ m/s}$, and $b := 0.275 \text{ m/s}$.

Combining all of the above ingredients, the optimization of pitch angles can be written as

$$\begin{aligned} & \text{maximize} && \sum_{i=1}^n p_i G(\beta_i^T x) \\ & \text{subject to} && x^{\min} \leq x \leq x^{\max}, \\ & && \sum_{j=1}^m \frac{|\gamma_j|}{ax_j + b} \leq T, \end{aligned}$$

where p_i is the importance level of region i ; G is a non-negative, strictly monotonically increasing, and concave function; β_i is the vector containing the entries β_{ij} defined in (10); x is the optimization vector defined in (11), with bounds x^{\min} and x^{\max} ; $|\gamma_j|$ is the length of segment γ_j ; a and b are the constants in (5); and $T > 0$ is the user-defined time budget.

For the experiments considered in this paper, the minimum and maximum pitch angles were 15° and 35° , respectively. This corresponds to horizontal speeds in the range of $v_{\text{hor},j} \in [0.240, 0.262] \text{ m/s}$. For these ranges of values, the optimization is convex. To see this, note that the objective function is concave. Thus we are maximizing a concave function (or equivalently, minimizing a convex function $-C(x)$). The constraints in $x^{\min} \leq x \leq x^{\max}$ form a convex set. Finally, since $x_j < b/a \approx 5.5$ for all $x_j \in [\tan 15^\circ, \tan 35^\circ]$, the denominator of the time constraint $ax_j + b$ is strictly positive, and the time constraint also yields a convex set. Therefore, the optimization can be efficiently solved using standard convex optimization tools (Boyd and Vandenberghe, 2004). In the rest of the paper we present experimental results combining the path planning and speed optimization for adaptive ocean sampling.

5.1 Combined Iterative Algorithm

The two algorithms described above are designed to be used in tandem, with the path output of the ZZTOPP algorithm to be used as the input for the pitch angle optimization. With some minor modifications, the optimal pitch angles can also be used as the input for the ZZTOPP path planner, thus suggesting an iterative procedure in which the output of one algorithm is repeatedly used as the input for the other. Specifically, the cost function $H(W, \lambda)$ can be redefined as

$$H(W, \phi, \lambda) := \lambda \sum_{j=1}^m \int_{\gamma_j} \nu(q) dq - (1 - \lambda) \sum_{i=1}^n p_i \sum_{j=1}^m \frac{l_{ij} \tan \phi_j}{2(d_j^{\max} - d_j^{\min})}.$$

Notice that the second term is the same as the pitch angle cost function $C(\phi_1, \dots, \phi_m)$ from (9) with the function $G(z) = z$, and the first term is independent of the pitch angles. Therefore $H(W, \phi, \lambda)$ can be a common cost function for both the path and the angles. Let the time for traversal of a path be given by $\tau(W, \phi) = \sum_{j=1}^m \frac{\gamma_j}{a \tan(\phi_j) + b}$. Then the iterative algorithm is given by Algorithm 2.

One important point deserves comment. Notice that the time constraint in the angle optimization is given not as the global time constraint T , but as the traversal time of the previously found optimal path $\tau(W^*(t), \phi^*(t-1))$. This is required for the iterative algorithm to converge to practically useful paths. To see why this is the case, consider using the global time constraint T for both optimizations. The ZZTOP algorithm, being a discrete optimization, finds a path that is always some finite amount of time faster than the time constraint T . Conversely, the angle optimization, being a continuous optimization, finds solutions which precisely meet the maximum time constraint. After one execution of the iteration, the small amount of slack in the time

Algorithm 2: Path-Angle Iteration

```
1 Set  $\phi^*(0) \leftarrow [\phi_1(0) \cdots \phi_m(0)]^T$ 
2 Set  $t \leftarrow 1$ 
3 while stopping conditions are not met do
4    $W^*(t) \leftarrow \arg \min_W H(W, \phi^*(t-1), \lambda)$  subject to  $\tau(W, \phi^*(t-1)) < T$ 
5    $\phi^*(t) \leftarrow \arg \min_\phi H(W^*(t), \phi, \lambda)$  subject to  $\tau(W^*(t), \phi) < \tau(W^*(t), \phi^*(t-1))$ 
6    $t \leftarrow t + 1$ 
```

constraint from the ZZTOP algorithm is taken up by the angle optimization, in the form of slightly increased angles. In the next iteration, the ZZTOP algorithm finds a path that is slightly shorter along the legs with small angles and slightly longer along those with high angles, again with some slack in the time constraint. The angle planner again takes up this slack by increasing the angles. Over many iterations, this results in excessively short paths, with all angles set to the maximum allowable angle. The simple way to prevent this accumulating error and correction mechanism is to remove the slack in the time constraint for the angle optimization. Thus the time of the path given by the ZZTOP algorithm becomes the time constraint for the angle optimization.

At the time that the experiments in this paper were conducted, this iterative algorithm was still under development. For both of the SCB experiments we used a single execution of the ZZTOP algorithm, followed by a single execution of the angle optimization. For the MB experiment we used the path and angles resulting from five iterations of the ZZTOP algorithm and angle optimizations, using the global time constraint for both. Beyond the fifth iteration, the paths began to suffer noticeably from the slack accumulation mechanism described above. However, we found that Algorithm 2 terminates in two iterations for the SCB environment in this paper, giving results very similar to the paths and angles used for the experiments. Likewise, for the MB environment Algorithm 2 terminates in four iterations with results very similar to the path and angles used for that experiment. In future experiments we expect Algorithm 2 to prove a useful means of obtaining good paths with minimal user intervention.

6 Ocean Science Applications

The work in this paper is motivated through practical applications in coastal marine science. We consider the general, age-old, oceanography problem of creating a long-term time-series of measurements for a specific area of the ocean, but also consider gathering data for the study of a specific biophysical process that has signatures that oscillate at multiple spatiotemporal resolutions. Specifically, we are interested in developing tools to help ocean scientists decipher the environmental triggers leading to the onset, evolution and eventual mortality of Harmful Algal Blooms (HABs). This motivational example persists in many coastal communities around the world, and in particular, southern California and Monterey Bay, CA.

We present experimental results of three field deployments that implement the algorithms presented in the previous sections. Two of the deployments occurred within the Southern California Bight (SCB), and one deployment occurred in Monterey Bay (MB). The first deployment in the SCB was the subject of preliminary work in this area by the authors (Smith et al., 2011). Here, we recount this initial study and extend our previous efforts through a deployment in MB and another study in the SCB.

Many processes and phenomena that we aim to study in the SCB, e.g., algal blooms, also occur in MB, although the triggers, residence time and forcing parameters are slightly different in each region. Our proposed algorithms are designed to produce a sampling method that acquires data for processes occurring at multiple spatial resolutions. In many cases, an appropriate sampling resolution is not precisely known. However, the general area of study, or region of interest, and a process or feature of interest are known. Given an oceanic region of interest, an ocean scientist with expert domain knowledge can identify locations

within the region that are of greater interest than other areas. For example, in shelf break regions, the water depth decreases rapidly toward shore. The interaction of physical forcing with the bathymetry and coastline can result in coastal upwelling, diapycnal mixing, and generation of internal tides and waves that may break near the coast, all of which contribute to upward transport of nutrients in the coastal region. Upwelling brings colder, nutrient-rich water into the euphotic zone, thus providing increased phytoplankton growth and production. This is one conjecture for the formation of algal blooms. The upwelling also brings denser water to the surface and creates an unstable equilibrium that may result in both horizontal and vertical mixing. Hence, in shelf-break regions, it is advantageous to collect higher resolution data. Thus, an ocean expert may be aware of *high-interest* areas within a region, but may still not fully understand the forcing, dynamics and/or biological response that make those areas interesting. These high-interest areas can exist at different temporal scales as well. For instance, a region like the shelf break may be of high interest regardless of the season, whereas the area near a river mouth may only be of high interest immediately following a rain event when large quantities of buoyant, nutrient-rich runoff are discharged at the coastal boundary. The ability for an ocean scientist to determine and rank areas of high interest within our mission planner ensures that the computed path visits high interest regions and allocates an appropriate sampling resolution based on the relative importance within the survey region.

In the following sections, we present a brief background on the two regions of study, the SCB and MB, and outline some of the factors to consider when planning sampling missions in each region.

6.1 The Southern California Bight

Our efforts in the Southern California Bight have been specifically focused on a region below the Palos Verdes peninsula that is commonly referred to as San Pedro Bay. This is a region where three major river systems discharge and is bounded by two large ocean outfalls that discharge about 600 million gallons of treated sewage per day. Coastal ocean processes in southern California are driven typically by large-scale processes, rather than local forcing, e.g., local winds (Jones et al., 2002). The dynamics are affected by multiple processes that are characterized by many different spatiotemporal scales that cannot be fully resolved by use of a few stationary sensors or via a single, short-term AUV deployment. A significant, long-term data set is required to understand large-scale variability in this complex coastal ecosystem. Smaller-scale processes also significantly impact the biological dynamics in the SCB. Three examples include river runoff into the ocean, sewage outfalls, and the propagation of internal waves. In San Pedro Bay, south of Los Angeles and Long Beach harbors, two major rivers inject storm water and tertiary treated effluent directly into the coastal surface waters. Additionally, the greater Los Angeles area discharges nearly one billion gallons per day of blended primary and secondary effluent into coastal waters via offshore outfall diffusers that are located nominally at about 60 meters depth, and within 10-12 km of the coast. In contrast to these two forms of anthropogenic input, coastal upwelling and internal waves are two natural phenomena that provide nutrient fluxes into the coastal ocean's upper layer. Internal waves that arise when coastal currents and tides interact with the coastal topography generating waves that propagate along the subsurface pycnocline (density gradient). These waves may break as water depth decreases across the shelf, or where the pycnocline intercepts the bottom, contributing to vertical mixing across the pycnocline.

Only from the comprehensive view that these long-term data sets provide can we then begin to isolate the important locations or dominant biogeochemical properties that drive the system. And, through multiple years of historical data comparison, we can begin to see the frequency and impact of long term climatological processes. For more information on the physical and biological dynamics in the San Pedro Bay and SCB, we refer the interested reader to (Schnitzer et al., 2007; Noble et al., 2009; Smith et al., 2010b).

6.2 Monterey Bay, California

Similar to the SCB, MB is also affected by large-scale processes, however, these are not the primary driving processes. Local wind forcing and unique topography make MB dynamics particularly interesting. The

Monterey Canyon, one of the largest underwater canyons in the world, extends from the coast at Moss Landing, CA, in the center of Monterey Bay, 153 km into the Pacific Ocean where it terminates at the Monterey submarine fan. This canyon reaches depth of up to 3,600 m. The bay itself is a productive coastal environment located within the central California Current System. Wind-driven coastal upwelling and Ekman pumping as well as internal tides provide nutrients to the surface layer that support high primary productivity. The water in MB is a mixture of colder, more saline upwelled water and warmer, fresher water from the California Current System. Based on the relative proportions of these sources, the environment in MB changes significantly. Hence, MB is primarily influenced by the oceanographic dynamics resulting from local and regional forcing, such as cycles of upwelling, favorable local winds, and reversals.

It has recently been shown that the waters of northeast MB function as an *extreme bloom incubator* (Ryan et al., 2008), frequently developing dense *redtide* algal blooms that can rapidly spread to other areas along the coast. The intensity and biological consequences of coastal upwelling in this region are greatest between March and November (Pennington and Chavez, 2000). For more information on the physical and biological dynamics of MB, we refer the interested reader to (Ryan et al., 2008; Ryan et al., 2009; Ryan et al., 2010).


7 Field Experiments in San Pedro Bay and Monterey Bay

For demonstration and validation of the algorithms and their outputs presented in the previous sections, we consider data collected from sea trials. We deployed a Slocum glider in San Pedro Bay within the SCB for three weeks on two separate occasions, and in MB for two weeks, for the implementation of our computed missions. For this study, we compared the execution of reference paths to those paths computed by use of our method. Since both the SCB and MB are under active investigation by expert researchers, see e.g., (Smith et al., 2010b) and (Pennington and Chavez, 2000; Ryan et al., 2009), the *reference* path is one designed by an ocean scientist with extensive domain knowledge. This path is traversed multiple times to provide a baseline for comparison. In the case of the SCB experiments, the reference path has been repeatedly executed over the last two years. All reference paths were executed using a standard operating procedure; constant dive and ascent pitch angles of 26° .


The experimental results are divided into four parts to test separate portions of the algorithms. First, we will compare collected science data in a region of high interest for the minimum, standard and maximum dive and ascent angles, 15° , 26° and 35° , respectively. Secondly, we compare the traversability of the two different paths, i.e., which path was followed more accurately by the glider. Thirdly, we compare the total time of traversal for the three missions. Lastly, we provide an overall assessment of the computed path and the implementation as compared with the reference path in the given region. In this study, we do not present a comprehensive analysis of the collected science data, as the metrics for evaluation chosen here are specifically aimed at the implementability and overall design of the compared paths. Additional, multi-vehicle field trials are required to adequately compare and contrast the data collection along each considered path.

7.1 Variable Resolution


Here, we examine the variability in the sample resolution that can be collected by the gliders. Based on experimental trials, a safe operational range for ascent and descent pitch angles for the glider was determined to be $15^\circ - 35^\circ$. To investigate the difference in sampling resolution for this range of pitch angles, we executed cross-shelf transects through region Q_1 (see Figure 8) at pitch angles of 15° , 26° and 35° . The path for these transects is γ_3 of the reference path for the SCB, for comparison at a later date with previously collected science data. For evaluation purposes, the temperature ($^\circ\text{C}$) data collected during these experiments is presented in Figs. 7(a), 7(b) and 7(c). In these figures, we display the individual measurements taken by the vehicle to emphasize the difference in sampling resolution. As previously mentioned in Section 5, we only gather data when the glider is descending, or on the downcast. The collection of samples taken on a single descent is called a profile. These individual depth profiles are generally interpolated to a standard resolution




(a) Temperature data with 15° pitch angle.



(b) Temperature data with 26° pitch angle.



(c) Temperature data with 35° pitch angle.



(d) Interpolation of the data in Fig. 7(c) over the entire transect.

Figure 7: Pressure vs. Distance from Shore plots of Temperature ($^\circ\text{C}$) data collected along the same transect for 7(a) 15° , 7(b) 26° and 7(c) 35° . Individual measurements are shown to emphasize the sampling resolution along each path. An interpolation of the data collected over the entire transect executed by the glider with a 35° pitch angle is given in Fig. 7(d). Start and end times for the transect are shown above each respective figure, the year is 2010, and times listed are GMT.

Table 1: Calculated and observed sample density for pitch angles of 15° , 26° and 35° for a depth range of 0 – 80 m.

	15°	26°	35°
Calculated Distance between profiles (m)	597	328	229
Calculated Profiles per km	1.67	3.05	4.38
Observed Profiles per km	2.6 ± 0.05	3 ± 0.08	4.1 ± 0.1

grid to generate a continuous spatiotemporal representation of the water column.

From the data presented in Figures 7(a), 7(b) and 7(c), it is clear that by increasing the pitch angle of the glider’s dive and ascent, we can alter the resolution at which the data are gathered. Based on the glider trajectory, the theoretical sampling density is given by Equation (8). For pitch angles of 15° , 26° and 35° , using Equation (8) to calculate the theoretical distance between these depth profiles (samples) and the associated profile (sample) density (per km traveled) gives the values in the first two rows of Table 1. From the data presented in Figures 7(a), 7(b) and 7(c), and two additional transects for each angle, we can calculate the average number of samples per kilometer traveled taken by the glider for each prescribed angle. Considering the region shown in Figure 7, we observe that theoretical predictions for the sampling densities in the depth range 0 – 80 m matches well with experimental results, see Table 1. In particular, we see that in the range of angles considered, a 10° increase in the pitch angle provides, on average, one additional depth profile (sample) per kilometer traversed. Given that we typically command the glider to surface every six hours, an average continuous transect is four kilometers. The difference between the same 4 km transect executed with pitch angles of 15° and 35° is 10 and 16 depth profiles, or samples, respectively. Thus, we achieve more than a 60% increase in spatial sampling resolution through a 20° change in the glider’s executed pitch angle. This is a significant enough increase to investigate further exploitation of this technique, especially in matching the sampling resolution with the specific oscillatory frequency of a given process.

Previous data collects have presented undulations of the thermocline that appear to have a scale of ~ 2 km. Based on satellite imagery of the region, we believe that that these undulations are due to larger-scale features advecting through the region. Previous surveys of this region with the glider, have shown some internal wave trains on the shelf with wave amplitudes of ~ 5 m, and with a wavelength of ~ 350 m. These wavelengths can be slightly compressed as the wave train progresses shoreward and the glider moves seaward. Thus, resolution of an internal wave is highly dependent on the spatial resolution of the glider profiles which are a function of the pitch (dive) angle, glider speed, and profile depth.

From the data presented, if an overall increase in spatial sampling resolution is desired, this can be achieved by prescribing that the glider operate with the steepest pitch angle for the entire mission. This would suffice for gathering data related to fine-scale processes. However, in the experiments presented in the sections to follow, we are also interested in gathering data for a large-scale algal bloom event. Based on a 10-day life-cycle of such an event, we constrain the total time to execute one loop of the computed path to five days. This constraint forces the glider to use pitch angles less than the maximum to satisfy the prescribed time budget.

7.2 Deployment 1: SCB

We present experimental results from a field deployment in the SCB from August 11, 2010 through September 8, 2010. For this implementation, the consideration for the ocean currents is based on the temporally averaged magnitudes from ROMS. Note that the set of points defining the ROMS daily output is denoted by M . We consider historical ROMS predictions for 30 days prior to the deployment. For each daily prediction, we consider 24 hours of the forecast. At each grid point in M , we find the maximum magnitude of the ocean current between 0 and 80 m depth. Since strong currents in any direction effect the glider’s navigational



Figure 8: A general overview of the experimental area off the coast of southern California. The mission domain Q is delineated by the white polygon. Regions Q_1 , Q_2 and Q_3 are the high interest regions, and are delineated by the red, green and yellow polygons, respectively. The grey shaded polygons denote the checkerboard regions, and the yellow regions to the west of the mission domain are the primary shipping lanes for Long Beach Harbor.

accuracy, we only consider magnitude, and do not take current direction into account at this phase of the study. The maximum magnitudes at each location for each day are then averaged over the 30-day time window. By interpolating between the grid points in H , we create a function $\nu : Q \rightarrow \mathbb{R}_{\geq 0}$, which gives the average maximum magnitude current expected at each point in Q .

7.2.1 Path Planning and Algorithm Details

For the application of Alg. 1, we chose six waypoints. In Figure 8 we display the region of interest Q , L , the primary shipping lanes that are to be avoided, high-interest regions Q_1 , Q_2 and Q_3 , and the six computed checkerboard regions. For this experiment, one cycle of γ must be traversable in < 5 days. Substituting 26° for ϕ_j in Equation (12), this corresponds to a total path length of 108.26 km, so we assume that γ must be less than 110 km. The other inputs to ZZTOPP for the path optimization are the user-defined parameters obtained from an ocean scientist with expert domain knowledge: $\lambda = 0.4$, $p_1 = 1$, $p_2 = 0.75$, $p_3 = 0.7$, the background importance $p_n = 0.3$, and the location of checkerboard region $j = 1$ is the northeast corner of Q . Here, λ is chosen such that the magnitude of the mean magnitude currents are the same order of magnitude as the region weights. The weights p_i were provided by an ocean scientist with expert domain knowledge. This mission was created by use of the ZZTOPP algorithm from Section 4, followed by an application of the pitch angle planning algorithm from Section 5.

For this experiment, we had to extend the checkerboard region $j = 5$ to include the area within Q_2 , see Figure 9. Based on the location of Q_2 and checkerboard region $j = 1$, odd m is operationally unsafe, as the path from w_m to w_1 could cross over land. For m even, the intersection of Q_2 and checkerboard region $m - 1$ was not a large enough area to guarantee a minimal path through Q_2 .

Applying the ZZTOPP algorithm with the given inputs produces the black path in Figure 9. The magenta path in Figure 9 is the reference path that was hand-designed by an ocean scientist for the same application presented here.




Figure 9: This figure presents an enlarged view of Figure 8 with the implemented paths overlaid. The path produced by the ZZTOPP algorithm is given by the black line, while the magenta line represents the reference path that was hand-delineated by an oceanographer to survey the same region with the same constraints. The mission domain Q is delineated by the white polygon. Regions Q_1 , Q_2 and Q_3 are the high interest regions, and are delineated by the red, green and yellow polygons, respectively. The high-interest region weights are $p_1 = 1$, $p_2 = 0.75$, $p_3 = 0.7$, with the background importance $p_n = 0.3$. The grey shaded polygons denote the checkerboard regions.

7.2.2 Navigational Accuracy


To assess the effectiveness of the ocean current consideration used in the ZZTOPP algorithm, we compare the navigational accuracy of the three implemented experiments. We compare the prescribed path with the executed path by use of the following metric. We delineate the glider's executed path by connecting sequential surfacing locations during the mission. We then compute the absolute area between the executed path and the prescribed path. This essentially integrates the positional error along the entire closed-loop path. This area measure is used as the navigation score, with a smaller score indicating a more accurately navigated path.

For the reference path with the standard implementation, we average the results from ten recent loop traversals. Experimental results are presented in the first column of Table 2, including one standard deviation uncertainties.


For these ten standard executions of the reference path, we see an average navigation score of 70.35 km^2 with a standard deviation of 13.35 km^2 . Figure 10(a) displays the best execution of the reference path with the standard pitch angles applied. For the execution of the reference path with pitch angles optimized for sample resolution, the navigation score increased to 86.06 km^2 . For the path computed by use of the ZZTOPP algorithm, implemented with optimized pitch angles, we get a navigation score of 56.23 km^2 . These two later paths with optimized pitch angles were only executed once during this experiment. Figures 10(b) and 10(c) display the path of the glider for the execution of the reference path with optimized pitch angles and the computed path, respectively. These results show that both paths fall within one standard deviation of the standard reference path executions. This motivates further trials of the proposed technique to verify that we can provide similar results to the reference path in navigational accuracy, while gathering data at a higher spatial resolution.

Table 2: Experiment statistics from the planned path compared with historical statistics from a reference path


	Reference Path (Standard)	Reference Path with Optimized Pitch Angles	Computed Path with Optimized Pitch Angles
Prescribed Path Length (km)	97.3	97.3	99.2
Pitch Angles ($\gamma_1, \dots, \gamma_6$)	($26^\circ, 26^\circ, 26^\circ, 26^\circ, 26^\circ, 26^\circ$)	($15^\circ, 27^\circ, 35^\circ, 35^\circ, 25^\circ, 15^\circ$)	($15^\circ, 35^\circ, 35^\circ, 25^\circ, 35^\circ, 35^\circ$)
Actual Distance Traveled (km)	93.51 ± 4.58	105	102
Total Traversal Time (hhh:mm)	$110 : 02 \pm 019 : 58$	126 : 26	115 : 53
Navigation Score (km ²)	70.35 ± 13.35	86.06	56.23
Navigation Score per km traveled (km)	0.76 ± 0.16	0.82	0.55
$H(W, \lambda)$	-20, 280	-20, 280	-24, 638



(a) The reference path for the SCB region (magenta line) and an example of the path executed by the glider (cyan line).



(b) The reference path for the SCB region (magenta line) and the execution of the reference path with optimized pitch angles (cyan line).



(c) The reference path for the SCB region (magenta line) and the execution of the computed path (cyan line).

Figure 10: Three planned and executed paths for the SCB region. Figure 10(a) gives the initial reference path, Figure 10(b) shows the reference path with optimized pitch angles, and Figure 10(c) displays the path computed by the ZZTOPP Algorithm.

Table 3: Planned and experimental statistical results for the first deployment in the SCB.

	Prescribed Path Length (km)	Theoretical # of Prescribed Profiles	Actual Path Length (km)	Executed Profiles
Q_1 - Reference	3.34	10	3.62 ± 0.7	11 ± 2.1
Q_1 - Reference (Optimized)	3.34	14	0.25	1
Q_1 - Computed	6.68	29	8.18	36
Q_2 - Reference	15.03	46	13.99 ± 2.04	49.9 ± 7.7
Q_2 - Reference (Optimized)	15.03	14	10.14	34
Q_2 - Computed	22.55	85	18.44	63
Q_3 - Reference	2.95	8	2.29 ± 0.48	8.5 ± 1.8
Q_3 - Reference (Optimized)	2.95	14	2.76	10
Q_3 - Computed	5.75	21	4.38	16

7.2.3 Loop Traversal Time

An important component of the path design is to assist in resolving the frequency of ocean phenomena occurring at different spatiotemporal resolutions. Thus, a computed path incorporating any variations in velocity must not be slower than the reference path with standard operational procedures. In Table 2 we present the total time of traversal for one loop of each of the three executed mission scenarios. For ten recent executions of the standard reference path, we see an average loop traversal time of 110.03 hours. This average lies well within the desired traversal time of 4–5 days (96–120 hours), as mentioned previously. For the reference path executed with velocity control, we have a loop traversal time of 126.43 hours. Although this time is greater than the 120 hours desired, it is not excessively long. The primary reason for the extra time required is the additional 9 km traveled during execution due to poor navigational accuracy. For the computed path executed with velocity control, we see a loop traversal time of 115.88 hours. This is slightly longer than the reference path with standard execution, but still lies well within the desired time range.

7.2.4 Path Comparison

The navigational accuracy and loop traversal time are good metrics to compare the implementation of the reference and computed paths. However, we would also like to assess the merit of the prescribed sampling resolution and how the path compares based upon the chosen optimization criteria. First, we remark that based on the optimization criteria, we see more than a 20% improvement in the path score, $H(W, \lambda)$, for the computed path over the reference path, see Table 2. Thus, we are effectively planning a path through areas of lower magnitude currents, as well as spending more time in areas of high interest. In Table 3, we give the length of the segments for the planned and executed paths that pass through regions of high interest, and the corresponding theoretical estimations of the number of profiles executed in these regions based on Equation (8), the prescribed pitch angle along the segment, and the data presented in Table 1. The number of profiles is rounded to the nearest integer value. We point out the significant increase in the total segment length, and hence the number of profiles taken in each of the high interest regions for the computed path as compared to the reference path in both the prescribed and executed scenarios. From the inputs to the algorithm, we note that region Q_1 was the most important. It is of interest to note that the reference path traverses this region only once, while the computed path crosses through this region twice in one closed-loop cycle. The computed path does sacrifice some sample coverage along the northeastern edge of the region of interest that is covered by the reference path. From Figure 9, we see that the computed path is skewed toward the western boundary of the region of interest because of the locations and weights assigned to high interest regions Q_1 and Q_2 .



Figure 11: The white polygon delineates the general survey Region Q_1 and the colored polygons designate the *high-interest* regions used as input to our mission planning algorithms. Region Q_1 (red) is the LATMIX region, Region Q_2 (yellow) is an area near an MBARI mooring buoy, Region Q_3 (cyan) was an area with an active algal bloom, and Region Q_4 (green) is an area near a river mouth.

7.3 Deployment 2: Monterey Bay

From October 7 - October 22, 2010, as part of the MBARI BIOSPACE program (Fulton-Bennett, 2010)⁵, the USC CINAPS glider was tasked with gathering data from the northeast region of Monterey Bay, CA. These data were then utilized during daily planning meetings to retask currently deployed vehicles and determine the utility of additional available assets. The result was a daily adaptation and replanning for the missions of all assets.

7.3.1 Path Planning and Algorithm Details


The initial, or *reference* trajectory for the Monterey Bay region was determined by biological oceanographer John Ryan, a Senior Research Specialist at MBARI. The mission objectives were to:

1. Provide a synoptic view of the northeast portion of Monterey Bay for depths > 25 m.
2. Execute cross-shelf transects from the 25 m isobath to the edge of the Monterey Canyon.
3. Revisit each location along the trajectory every 72 hours, i.e., 3-day cycles.

The general MB area is shown in Figure 11, with Q_1 , Q_2 , Q_3 and Q_4 denoting the designated regions of high-interest.

The initially delineated reference path (Fig. 12(a)) was executed from October 7 - October 13, 2010. On October 14, 2010, an adaptation to this reference path (Fig. 12(a)) was prescribed. The minor alteration was the addition of one waypoint in the northernmost corner of the original reference path. This new waypoint

⁵The BIOSPACE program is a multi-institutional collaboration funded by the U.S. Navy that constitutes one piece of the overarching MBARI CANON (Controlled, Agile, and Novel Observing Network) project (Chavez, 2010).



(a) Initial delineation of the reference path (magenta line) for the Monterey Bay experiment.

(b) Altered reference path (magenta line) for the Monterey Bay experiment.

Figure 12: Two reference paths executed during the BIOSPACE experiment in Monterey Bay. Figure 12(a) gives the initial reference path and Figure 12(b) shows the altered reference path.



(a) Initial reference path (magenta line) for the Monterey Bay experiment, and an execution (orange line) of this path.

(b) Altered reference path (magenta line) for the Monterey Bay experiment, and an execution (orange line) of this path.

(c) Altered reference path (magenta line) for the Monterey Bay experiment, and an execution (orange line) of this path.

Figure 13: Three executions of the designed reference paths for the BIOSPACE experiment in Monterey Bay. Figure 13(a) gives the initial reference path, and Figures 13(b) and 13(c) show the altered reference path.

was only added to facilitate data transfer by use of Freewave radio modems, as part of a cost-effective, coastal, communication infrastructure as presented in (Pereira et al., 2009; Smith et al., 2009; Smith et al., 2010b). This updated reference path was prescribed from October 14 through October 18, 2010.

During the initial 11 days of deployment, three closed-loop circuits were completed along the prescribed reference paths. Three distinct loops that followed the respective prescribed reference paths are shown in Figure 13. Specific details regarding the implementation of these reference paths are presented in Table 4.

After 11 days of data collection, we implemented the mission planning techniques previously presented to create a path for the glider that satisfied the same criteria as the reference path, but also traversed through regions of particular scientific interest. With the guidance of MBARI oceanographers, we were able to delineate and rank four *high-interest* regions for input to our proposed algorithm. In MB, there are two areas of the bay that are consistently scientifically interesting. These include areas surrounding the MBARI moorings and the northeast corner of the bay where the ONR DRI Scalable Lateral Mixing and Coherent Turbulence project, or LATMIX project is currently ongoing. Driven by the recent data collections, there were an additional two areas within the general survey region which were also identified as high-interest. All four regions are shown graphically in Figure 11. Region Q_1 and Region Q_2 are areas that are interesting




Figure 14: The white polygon delineates the general region of interest, with the colored polygons (red (Q_1), yellow (Q_2), cyan (Q_3) and green (Q_4)) depicting the designated *high-interest* regions. The original reference path is given by the magenta line and the optimal path computed by use of our iterative algorithm is given by the black line.

throughout the year, and provide data relating to the overall bay dynamics. Region Q_1 is the LATMIX Region and region Q_2 is an area surrounding an MBARI mooring buoy. Regions Q_3 and Q_4 are regions that were selected based on recent data collects and current weather conditions. Region Q_3 was particularly of interest since other assets determined that there was an algal bloom patch residing in that general area. Region Q_4 was highlighted based on the occurrence of rain event on October 16 and 17, 2010 in Moss Landing, CA. This region is near a river mouth, and is significant to determine anthropogenic nutrient inputs from river runoff.

The weights given to Regions $Q_1 - Q_4$ and the background region were 0.9, 1, 0.9, 0.8 and 0.45, respectively. These weights p_i were provided by MBARI ocean scientists who have an intimate understanding of the survey area. The path computed by the ZZTOPP and angle optimization algorithm, as described in Algorithm 2, is presented in Figure 15. The starting point was chosen to be a location in the northwest based on the estimated location of the next vehicle surfacing, and the direction of execution was chosen to be counter-clockwise so as not to fight the currents imposed by the existing cyclonic eddy in the region. The optimized angles along the segments of the closed-loop path are 31.2° , 32.8° , 32.9° , 35° , 15° , 23.2° .

Based on time constraints and remaining, on-board power, we were only able to execute a portion of one cycle of the path planned by use of the techniques presented in this paper. In the following two sections, we compare the execution of this path to the execution of the five cycles along the initial reference path. As in the previous sections, the analysis of the experiment is done by examining two metrics; the navigational accuracy and the traversal time.

7.3.2 Navigational Accuracy

For the planning of this path, the consideration for ocean currents utilized the NCOM ICON ocean model run by NRL. Three weeks of outputs were used to calculate the speed variance as presented in Section 3. We remark that the use of this method provides information about the variance of the ocean current in the region, but not the magnitude, as considered in the previous section. We assume that areas with low variance will be easier to predict, thus we can compensate for the currents in these regions. At this stage of




Figure 15: The white polygon delineates the general region of interest, with the colored polygons (red (Q_1), yellow (Q_2), cyan (Q_3) and green (Q_4)) depicting the designated *high-interest* regions. The computed path is given by the magenta line and the execution of this path is given by the orange line.

Table 4: Experimental statistics from the *long* reference paths.

	Long Circuit 1	Long Circuit 2	Long Circuit 3
Prescribed Path Length (km)	66.9	72	72
Actual Distance Traveled (km)	71	90.8	42.6
Pitch Angles ($\gamma_1, \dots, \gamma_6$)	($26^\circ, 26^\circ, 26^\circ, 26^\circ, 26^\circ, 26^\circ$)	($26^\circ, 26^\circ, 26^\circ, 26^\circ, 26^\circ, 26^\circ$)	($26^\circ, 26^\circ, 26^\circ, 26^\circ, 26^\circ, 26^\circ$)
Total Traversal Time (hh:mm)	61 : 53	98 : 05	39 : 07
Navigation Score (km^2)	64.2	174.3	29.9
Navigation Score per km traveled (km)	0.9	1.9	0.7
$H(W, \lambda)$	-117.65	-127.24	-127.24

the work, the compensation for the currents is not included. However, this is an area of ongoing work, with plans to incorporate the methods of increasing navigational accuracy by use of ocean models presented in (Smith et al., 2010d). For this reason, analyzing navigational accuracy is not particularly enlightening since we are not avoiding areas of predicted high magnitude currents, but only of high temporally varying currents. Additionally, the currents experienced in Monterey Bay are greater and more dynamic than those observed in the SCB, see Figures 4 and 5. In particular, for the survey region selected for our vehicle, a cyclonic eddy existed for the duration of the deployment. Such events have a detrimental effect on the navigational accuracy

7.3.3 Loop Traversal Time

The planned loop traversal time for this path was set to be ~ 70 hours. This planned time of execution for the big cycle corresponds to an average velocity of ~ 1 km /hr. Due to the sparse data set gathered and small number of complete cycles of any of the paths, we cannot perform an in-depth analysis for this metric of this experiment. We remark that we did not see any one path or mission being executed significantly

Table 5: Experiment statistics for the implementation of the computed path.

	Optimal Path
Prescribed Path Length (km)	70
Actual Distance Traveled (km)	48.2
Pitch Angles ($\gamma_1, \dots, \gamma_6$)	(31.2°, 32.8°, 32.9°, 35°, 15°, 23.2°)
Total Traversal Time (hh:mm)	54 : 52
Navigation Score (km ²)	41.2
Navigation Score per km traveled (km)	0.86
$H(W, \lambda)$	-191.6793

Table 6: Planned and experimental statistical results for the deployment in the Monterey Bay.

	Prescribed Path Length in Region (km)	Theoretical # of Prescribed Profiles	Actual Path Length in Region (km)	Number of Executed Profiles
<i>A</i> - Reference 1	3.64	11	5.01	15
<i>A</i> - Reference 2	9.1	28	6.02	18
<i>A</i> - Computed	12.66	46	10.57	37
<i>B</i> - Reference 1	5.75	18	0	0
<i>B</i> - Reference 2	5.47	17	4.01	12
<i>B</i> - Computed	8	35	4.03	18
<i>C</i> - Reference 1	3.82	12	3	9
<i>C</i> - Reference 2	3.82	12	2.12	7
<i>C</i> - Computed	6.33	28	5.28	23
<i>D</i> - Reference 1	0	0	0	0
<i>D</i> - Reference 2	0	0	2.41	7
<i>D</i> - Computed	0.1	0	2.1	4

faster or slower than any other. However, we do note that due to the larger magnitude and cyclonic currents experienced in MB, as compared to those in the SCB, it is important to attempt corrections for currents on vehicles like autonomous gliders to keep them from grossly navigating off course. Such an event can be seen in the data for the execution of Long Circuit 2 in Table 4. Here, a strong cyclonic eddy, coupled with the prescribed path, made it difficult for the glider to compensate for the ocean currents by use of the on-board algorithm.

7.3.4 Path Comparison

Based on the optimization criteria, we see a 63% improvement in the path score, $H(W, \lambda)$, for the computed path over the first reference path, and a 51% improvement for the computed path over the second reference path, see Table 5. Note that since we input the speed variance of the current, rather than the current magnitude as in the previous deployment, there is a difference in the order of magnitude of $H(W, \lambda)$. This is a region specific path score and is not intended to be viewed across different regions and deployments, especially since the underlying inputs to the algorithm are different. As seen in the previous SCB deployment, we are able to plan a path that meets preset science goals, and spends increased time in areas of high interest. As seen previously, there is a significant increase in the total segment length through the regions of interest for the computed path as compared to both reference paths. For regions Q_1, Q_2 and Q_3 , we see more than a 40% increase in path length inside regions of high interest in all but one case. Additionally, since the time constraint for completing one circuit was sufficiently large, our algorithm was able to assign pitch angles of nearly 35° to all but two segments. This coupled with the increased path length through the defined regions directly correlates to the significant increase in the number of profiles taken in each of the high-interest




Figure 16: The white polygon delineates the general region of interest, with the colored (red, yellow and green) polygons depicting the designated *high-interest* regions. The original reference path is given by the magenta line and the optimal path computed by use of our algorithm using current speed variance is given by the cyan line.

regions. The most significant increase is the difference between the length of path in Region Q_1 between the first delineated reference path and the computed path; more than a 300% increase in path length and more than a 400% increase in the number of profiles.

7.4 Deployment 3: SCB

In this section, we present the results of an additional field trial performed in the SCB. From November 4, 2010 to November 22, 2010 we implemented a newly designed sampling mission created by use of the ZZTOPP algorithm followed by the pitch angle optimization algorithm. For this experiment in the SCB, we consider the same general survey area and regions of high-interest as presented in Section 7.2, and shown in Figure 8. Additionally, we compare the results to the same reference path described in Section 7.2, and given by the magenta line in Figure 9. However, instead of incorporating ocean model outputs via consideration of large magnitude currents, we examine the ocean current variability by use of the analysis presented in Section 3. As mentioned in the previous section, this method of consideration will not tend towards keeping the glider out of areas of large magnitude currents, but will steer the glider through areas where the predictability of currents is better. We provide an assessment comparing the two methods of ocean current consideration, i.e., avoiding large magnitudes versus high variability.

7.4.1 Path Planning and Algorithm Details

The planned path for this experiment is different from that given in Section 7.2, as this path was designed with the current variance rather than the current magnitude for the ZZTOPP cost function. The output from the algorithm yields the path presented in Figure 16. Comparing the cyan path shown in Figure 16 to the magenta path shown in Figure 9, we see a very similar plan. This is expected because the checkerboard regions that define the potential successive waypoints are the same for both paths. However, we do notice two key differences between the two computed paths. The first difference is in the computed path. The path

Table 7: Experiment statistics from execution of the computed path, compared with historical statistics from the reference path from previous deployments.


	Reference Path (Standard)	Computed Path with Velocity Control 1	Computed Path with Velocity Control 2
Prescribed Path Length (km)	97.3	89.8	89.8
Pitch Angles ($\gamma_1, \dots, \gamma_6$)	($26^\circ, 26^\circ, 26^\circ,$ $26^\circ, 26^\circ, 26^\circ$)	($18^\circ, 35^\circ, 35^\circ,$ $15^\circ, 15^\circ, 35^\circ$)	($18^\circ, 35^\circ, 35^\circ,$ $15^\circ, 15^\circ, 35^\circ$)
Actual Distance Traveled (km)	93.51 ± 4.58	91	99.5
Total Traversal Time (hhh:mm)	110 : 02 \pm 019 : 58	105 : 07	132 : 54
Navigation Score (km ²)	70.35 ± 13.35	65.9	113.74
Navigation Score per km traveled (km)	0.76 ± 0.16	0.72	1.14
$H(W, \lambda)$	-93.31	-153.3334	-153.3334

presented in Figure 9 has better coverage in the southern portion of the survey region, however it avoids passing through region Q_1 on the segment connecting the southernmost waypoint with the northernmost waypoint. Conversely, the path given by the cyan line in Figure 16 has less coverage in the southern portion of the survey region, but has three segments cross through both Q_1 and Q_2 . In the later case, we see that the reward of the high interest regions dominated the cost of navigating along long path segments. Additionally, referring to Figure 5, we see that the computed path has less coverage in the southern portion of the survey region in an attempt to avoid higher variability current regimes. Since the computed variability for the survey region is relatively homogeneous, the increase in southwest corner drives the ZZTOPP algorithm to avoid this area. The second difference between the two experiments is in the sampling resolution, i.e., pitch angles along each segment, that is computed. We notice a different distribution of pitch angles for the path segments. This is a direct result of the two paths navigating through the high-interest regions in a different manner. The path computed for this experiment spends more time in high-interest regions than the computed path given in Section 7.2.

7.4.2 Navigational Accuracy

Reiterating, for ten standard executions of the reference path, we see an average navigation score of 70.35 km² with a standard deviation of 13.35 km². For the path computed by use of the iterative algorithm, we get navigation scores of 65.9 and 113.74 km² for the two executions of the computed path. To present a graphical representation of these results, we direct the reader to Figure 17. In Figure 17(a), we display the reference path with the magenta line and a typical execution of this path is represented by the cyan line. We remark that the executed paths are assumed to be the straight lines connecting the sequential locations at which the glider surfaced during the prescribed experiment. In Figure 17(b) we display the prescribed computed path for this deployment in magenta, with the path that the glider implemented for the second execution shown in cyan.

For the two executions of the computed path for this experiment, we see two distinct results. The first execution provides a promising result, with a navigation score (0.72) similar to the averaged navigation score (0.76) of the multiple executions of the reference path. The second execution falls short of expectations, yielding a navigation score of (1.14); almost three standard deviations from the averaged reference path results. Towards the end of the first execution, the ocean currents increased in magnitude, and maintained a steady NNW direction. These currents caused significant difficulties in navigation for the glider, as is seen in the data. In fact, a third execution of the computed path was attempted, however the magnitude and direction of the currents prohibited the glider from reaching the waypoint in the southeast corner of the survey region. Thus, we had to abort this trial.



(a) The reference path for the SCB region (magenta line) and an example path followed by the glider (cyan line).

(b) The computed path for the second SCB deployment (magenta line) and the path followed by the glider during the second execution (cyan line).

Figure 17: Prescribed (magenta) and executed (cyan) paths for deployments in the SCB region. High-interest regions Q_1 , Q_2 and Q_3 are given by the colored polygons; red, yellow and green, respectively, with the survey area delineated by the white polygon. The scale is the same in both figures.

7.4.3 Loop Traversal Time

In Table 7 we present the total time of traversal for one loop of each of the two mission executions, as well as ten recent executions of the reference path. For the reference path, we see an average loop traversal time of 110.03 hours. This is within the desired circuit time of 96 – 120 hours. For the computed path executions, we see a loop traversal time of 105.12 hours and 133 hours. The first execution agrees with the prescribed duration. The second execution exceeds the 120 hour upper bound by more than 12 hours, but still falls within one standard deviation of the average time of one execution of the reference path. This excess traversal time can be attributed to the glider having traveled an extra ~ 10 km. But, the primary reason for the large errors in navigational accuracy and traversal time is a result of the strong currents experienced during this experiment.

7.4.4 Path Comparison

This second deployment for the SCB not only gives us information regarding the path coverage and score compared to the heavily executed reference path, but it also provides some information on the inputs to the algorithms. The SCB deployment presented in Section 7.2 utilized a mean magnitude of the ocean currents to steer the glider through *calmer* waters. As previously mentioned, this deployment relied on the variability of the ocean currents to plan the path. From the three trials between the two deployments that implemented our computed path, we see that considering the average magnitude of the ocean current may be the better option. However, more deployments need to occur to validate this hypothesis. Additionally, to fully test the use of ocean variability, a trial needs to be carried out that utilizes ocean model predictions to assist the glider in compensating for currents. Coupling the path planning algorithms here with the navigational tools derived from ocean models presented in (Smith et al., 2010a) is an area of future work.

Now, we consider the computed path for this deployment. Based on the optimization criteria, we see more than a 64% improvement in the path score, $H(W, \lambda)$, for the computed path over the reference path, see Table 7. The path scores are comparable since the input for the ocean current consideration is the same as

Table 8: Planned and experimental statistical results for the second deployment in the SCB.

	Prescribed Path Length in Region (km)	Theoretical # of Prescribed Profiles	Actual Path Length in Region (km)	Number of Executed Profiles
Q_1 - Reference	3.34	10	3.62 ± 0.7	11 ± 2.1
Q_1 - Computed 1	10.03	44	8.98	31
Q_1 - Computed 2	10.03	44	6.85	22
Q_2 - Reference	15.03	46	13.99 ± 2.04	49.9 ± 7.7
Q_2 - Computed 1	20.9	81	20.89	92
Q_2 - Computed 2	20.9	81	22.89	100
Q_3 - Reference	1.95	5	2.29 ± 0.48	8.5 ± 1.8
Q_3 - Computed 1	3.16	5.3	2.84	5
Q_3 - Computed 2	3.16	5	3.05	5

that in Section 7.3. However, based on the fact that the time constraints and the high-interest regions are different, a direct comparison cannot be made. In Table 8, we give the length of the segments for the planned and executed paths that pass through regions of high interest, and the corresponding theoretical estimates of the number of profiles executed in these regions based on Equation (8), the prescribed pitch angle along the segment, and the data presented in Table 1. The number of profiles is rounded to the nearest integer value. As pointed out for the previous two deployments, we see an overall increase in the total segment length, and hence the number of profiles taken, in each of the high interest regions for the computed path as compared to the reference path. Comparing the computed path for this deployment with the computed path for the deployment in Section 7.2, we see a similar shift away from the eastern side of the survey s noted earlier. The coverage of this computed path additionally lacks in the southwest corner of the survey region as compared to the computed path of Section 7.2, however this computed path traverses through region Q_1 three times. The angles prescribed along each segment of the computed path are similar to those for the computed path in Section 7.2, given the areas that the segments traverse. It is of interest to note that this computed path is expected to gather $> 50\%$ more profiles than the computed path from Section 7.2, and $> 400\%$ more profiles than the reference path in region Q_1 ; the most important region.

8 Conclusions and Future Work

Our presented work proposed two algorithms that together produce paths for underwater gliders to provide persistent multi-scale resolution sampling. First, we computed a closed path to be continually traversed by a vehicle. Then, along this path we optimized the glider’s pitch angle to tune the spatial sampling resolution throughout the region of interest. These two algorithms were applied in tandem, with the output of the path planning algorithm being used as the input for the angle optimization algorithm. We also presented an iterative procedure in which Algorithm 1 and the angle optimization repeatedly use the counterpart’s output.

The missions computed by use of our techniques were implemented on autonomous gliders during three separate deployments; two in the Southern California Bight and one in Monterey Bay. In all cases, data collected from these experiments were compared with historical data from the execution of a reference path. We presented data from sea-trials that showed the computed paths covered more distance in less time, provided denser sampling within designated high-interest regions, and achieved a similar navigational accuracy to previous implementations of the reference paths. In addition, the computed paths also satisfied the same mission goals as the reference paths for the given region. The experimental results suggest that our algorithms provide the ability to plan missions for long-term, persistent monitoring to capture large-scale event frequencies, while additionally resolving smaller-scale events by locally modifying spatiotemporal sampling resolution.

The presentation here extends the preliminary results in (Smith et al., 2011) by performing additional field trials, deploying vehicles in a different geographical region, and providing two separate considerations for the incorporation of ocean currents in path planning for long-term monitoring deployments. These extended experiments further reinforce the findings of the initial study, however future deployments are necessary to examine long-term implementability issues, such as path repeatability, merit of collected science data, and whether the optimization parameters chosen here satisfy all the data collection parameters on a longer temporal scale.

The field experiments presented in this study demonstrates multiple key factors promoting our proposed method and the applications of robot-assisted ocean sampling. First, we have shown three distinct implementations of our mission planning algorithms to gather spatiotemporally variable data in a given region. Second, by experimenting in two separate regions, it is clear that we are not regionally constrained by any assumptions, and yet these assumptions are specific enough to gather data for multiple ocean science applications. A primary component that allows our methods to permeate multiple application areas is that we allow the user, e.g., ocean expert, to designate the closed-loop cycle time, and the relatively ranked regions of interest, but rely on no *a priori* knowledge of the underlying science to compute the mission. This allows the experience of the expert with extensive domain knowledge to incorporate their knowledge without specifically writing out complex cost functions and optimization parameters. Additionally, ocean science is a study that is particularly dominated by sampling scales, i.e., based on the process being studied one must consider the appropriate spatial and temporal scale at which to sample to resolve the feature. Here, this is encoded into our method through the user inputs. Third, the presented method is designed to incorporate vehicle constraints, e.g., operational velocity range, navigational accuracy, sampling method, and compute paths accordingly that are executable by the given platform.

For future study, we are planning sea trials in early 2011 to extensively test our computed mission plans for the SCB region. The primary goal of these trials is to compare and contrast the science data collected along the reference path versus along our computed path. This is an important assessment to properly understanding the scientific utility of the proposed technique for path planning and data collection. For this experiment, we plan to simultaneously operate two to three vehicles that are each executing a separate mission within the general area of interest. Hence, all vehicles will be subject to the same environmental effects, and be privy to collect data relating to the same biophysical process. Under these conditions, we will be able to provide an accurate assessment of the data collection abilities of our computed missions as compared to the reference path. We plan to review collected science data with an oceanographer with detailed domain knowledge to determine whether or not the sampling techniques presented here have provided *better* data with which to resolve small-scale events spatially while also collecting data related to large-scale processes in the region.

The collected science data from the deployments are a valuable motivation and outcome for this study. However, as with any field deployment, solving the asset allocation problem of having the vehicle, and hence the sensors, *in the right place at the right time* is an equally challenging and important dual problem to that of physically gathering the science data. Addressing this problem is also an area of future work in line with this research. Here, we have proposed the use of regional model outputs to design paths that steer a vehicle through regions of high interest, while also considering current magnitude and variability. The consideration for ocean currents is intended to operate the vehicle in areas where the external forces are small in magnitude and/or variability. Thus, compensation for these forces in the mission planning stages can be done. Additionally, in areas where ocean currents can be predicted well, we may seek to utilize string magnitude currents to increase glider speeds where necessary, or slow the velocity without altering the pitch angle. However, these ideas rely heavily on the ability to predict the ocean currents, and would require an advanced understanding of the region of interest. Combining the paths computed in this study with the techniques presented in (Smith et al., 2010a; Smith et al., 2010d) that increase navigational accuracy of autonomous gliders, will be the next phase of this study. This combined approach will hopefully lead to the design and implementation of long-term monitoring paths that can be repeatedly traversed with a reasonable accuracy.

9 Acknowledgments

The authors would like to thank Arvind Pereira, Matthew Ragan, Carl Oberg and Bridget Seegers for their valuable assistance with glider deployment, general operations and data processing. We would like to thank John Ryan and Sergey Frolov for their enlightening discussions and valuable information regarding the biophysical dynamics of Monterey Bay. We additionally thank MBARI for the opportunity to participate in BIOSPACE 2010 (Fulton-Bennett, 2010), a program funded by the U.S. Navy and carried out by the Monterey Bay Aquarium Research Institute (MBARI).

This work was supported in part by the NOAA MERHAB program (NA05NOS4781228), NSF as part of the Center for Embedded Network Sensing (CENS) (CCR-0120778), NSF grants CNS-0520305 and CNS-0540420, the ONR MURI program (N00014-09-1-1031 and N00014-08-1-0693), the ONR SoA program, and a gift from the Okawa Foundation. Funding for B. Jones comes from the USC Sea Grant Program, NOAA IOOS through the Southern California Coastal Ocean Observing System, and the California Ocean Protection Council through the Bight 10 program. The ROMS predictions used were computed by the Jet Propulsion Laboratory (JPL), California Institute of Technology, under a contract with the National Aeronautics and Space Administration (NASA). The NCOM ICON predictions used were computed by the United States Naval Research Laboratory (NRL) field site in Monterey, CA.

References

- Allmendinger, E. E. (1990). *Submersible Vehicle Systems Design*. The Society of Naval Architects and Marine Engineers.
- Anderson, D., Hoagland, P., Kaoru, Y., and White, A. (2000). Economic impacts from harmful algal blooms (HABs) in the United States. Technical report, Woods Hole Oceanographic Institution Technical Report: WHOI 2000-11.
- Anstee, S. and Fang, C. (2010). Coverage path planning for harbour surveys using an AUV. In *Proceedings of the IEEE Oceans Conference*, Sydney, Australia.
- Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.
- Bratkovich, A. (1985). Aspects of the tidal variability observed on the southern-california continental-shelf. *Journal of Physical Oceanography*, 15(3):225–239.
- Cetinic, I., Jones, B. H., Moline, M. A., and Schofield, O. (2010). Resolving urban plumes using autonomous gliders in the coastal ocean. *Continental Shelf Research*. In review.
- Chavez, F. (2010). CANON: Controlled, Agile, and Novel Observing Network. <http://www.mbari.org/canon/>. Viewed November 2010.
- Chelton, D. B., deSzoeke, R. A., Schlax, M. G., Naggar, K. E., and Siwertz, N. (1998). Geographical variability of the first baroclinic rossby radius of deformation. *Journal of Physical Oceanography*, 28(3):433–460.
- Choset, H. (2000). Coverage of known spaces: The boustrophedon cellular decomposition. *Autonomous Robots*, 9:247 – 253.
- Corcoran, A. A., Reifel, K. M., Jones, B. H., and Shipe, R. F. (2010). Spatiotemporal development of physical, chemical, and biological characteristics of stormwater plumes in santa monica bay, california (usa). *Journal of Sea Research*, 63(2):129–142. 569BI Times Cited:0 Cited References Count:85.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. (2001). *Introduction to Algorithms*. MIT Press, 2 edition.

- DiGiacomo, P. M., Washburn, L., Holt, B., and Jones, B. H. (2004). Coastal pollution-hazards in southern california observed by sar imagery: stormwater plumes, wastewater plumes, and natural hydrocarbon seeps. *Marine Pollution Bulletin*, 49(11-12):1013–1024. 882GH Times Cited:3 Cited References Count:62.
- Fossen, T. I. (1994). *Guidance and Control of Ocean Vehicles*. John Wiley & Sons.
- Fulton-Bennett, K. (2010). Multi-institutional experiment finds harmful algal blooms in Monterey Bay. http://www.mbari.org/news/news_releases/2010/canon/canon.html. Viewed November 2010.
- Graver, J. G. (2005). *Underwater Gliders: Dynamics, Control and Design*. PhD thesis, Princeton University.
- Griffiths, G., Jones, C., Ferguson, J., and Bose, N. (2007). Undersea gliders. *Journal of Ocean Technology*, 2(2):64–75.
- Hamilton, P., Noble, M. A., Largier, J., Rosenfeldd, L. K., and Robertstone, G. (2006). Cross-shelf subtidal variability in san pedro bay during summer. *Continental Shelf Research*, 26(6):681–702.
- Hoagland, P. and Scatasta, S. (2006). *The economic effects of harmful algal blooms*. Springer-Verlag.
- Jones, B. H., Noble, M. A., and Dickey, T. D. (2002). Hydrographic and particle distributions over the palos verdes continental shelf: spatial, seasonal and daily variability. *Continental Shelf Research*, 22(6-7):945–965.
- Jüttner, A., Szviatovszki, B., Mécs, I., and Rajkó, Z. (2001). Lagrange relaxation based method for the QoS routing problem. In *IEEE Conference on Computer Communications*, volume 2, pages 859–868, Anchorage, AK.
- Leonard, N. and Graver, J. (2001). Model-based feedback control of autonomous underwater gliders. *IEEE Journal of Oceanic Engineering, Special Issue on Autonomous Ocean-Sampling Networks*, 26(4):633–645.
- Leonard, N. E., Paley, D. A., Davis, R. E., Fratantoni, D. M., Lekien, F., and Zhang, F. (2010). Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in monterey bay. *Journal of Field Robotics*, 27(6):718–740.
- Li, Z., Chao, Y., McWilliams, J. C., and Ide, K. (2008). A three-dimensional variational data assimilation scheme for the regional ocean modeling system. *Journal of Atmospheric and Oceanic Technology*, 25:2074–2090.
- McGann, C., Py, F., Rajan, K., Ryan, J. P., and Henthorn, R. (2008). Adaptive control for autonomous underwater vehicles. In *Proceedings of the Association for the Advancement of Artificial Intelligence*, Chicago, IL.
- Noble, M., Jones, B., Hamilton, P., Xu, J., Robertson, G., Rosenfeld, L., and Largier, J. (2009). Cross-shelf transport into nearshore waters due to shoaling internal tides in San Pedro Bay, CA. *Continental Shelf Research*, 29(5):1821–1835.
- Oksanen, T. (2007). *Path Planning Algorithms for Agriculture Field Machines*. PhD thesis, Helsinki University of Technology, Espoo, Finland.
- Paley, D., Zhang, F., and Leonard, N. (2008). Cooperative control for ocean sampling: The glider coordinated control system. *IEEE Transactions on Control Systems Technology*, 16(4):735–744.
- Pennington, J. and Chavez, F. P. (2000). Seasonal fluctuations of temperature, salinity, nitrate, chlorophyll and primary production at station H3/M1 over 1989-1996 in Monterey Bay, California. *Deep-Sea Research II*, 47:947–973.
- Pereira, A., Heidarsson, H., Caron, D., Jones, B., and Sukhatme, G. (2009). An implementation of a communication framework for the cost-effective operation of slocum gliders in coastal regions. In *Proceedings of The 7th International Conference on Field and Service Robotics*, Cambridge, MA.

- Pitcher, G., Moita, T., Trainer, V., Kudela, R., Figueiras, P., and Probyn, T., editors (2005). *GEOHAB: Global Ecology and Oceanography of Harmful Algal Blooms, GEOHAB Core Research Project: HABs in Upwelling Systems*. SCOR and IOC, Baltimore and Paris.
- Ramsdell, J., Anderson, D., and Glibert, P., editors (2005). *HARRNESS: Harmful Algal Research and Response: A National Environmental Science Strategy 2005-2015*. Ecological Society of America, Washington, DC.
- Ryan, J. P., Fischer, A. M., Kudela, R. M., Gower, J. F., King, S. A., Marin, and Chavez, F. P. (2009). Influences of upwelling and downwelling winds on red tide bloom dynamics in Monterey Bay, California. *Continental Shelf Research*, 29(5-6):785–795.
- Ryan, J. P., Gower, J., King, S., Bissett, W., Fischer, A., Kudela, R., Kolber, Z., Mazzillo, F., Rienecker, E., and Chavez, F. (2008). A coastal ocean extreme bloom incubator. *Geophysical Research Letters*, 35.
- Ryan, J. P., McManus, M., and Sullivan, J. (2010). Interacting physical, chemical and biological forcing of phytoplankton thin-layer variability in Monterey Bay, California. *Continental Shelf Research*, 30:7–16.
- Schnetzer, A., Miller, P. E., Schaffner, R. A., Stauffer, B., Jones, B. H., Weisberg, S. B., DiGiacomo, P. M., Berelson, W. M., and Caron, D. A. (2007). Blooms of pseudo-nitzschia and domoic acid in the san pedro channel and los angeles harbor areas of the southern california bight, 2003-2004. *Harmful Algae*, 6(3):372–387.
- Schofield, O., Kohut, J., Aragon, D., Creed, E., Graver, J., Haldman, C., Kerfoot, J., Roarty, H., Jones, C., Webb, D., and Glenn, S. (2007). Slocum gliders: Robust and ready. *Journal of Field Robotics*, 24(6):473–485.
- Sekula-Wood, E., Schnetzer, A., Benitez-Nelson, C., Anderson, C., Berelson, W., Brzezinski, M., Burns, J., Caron, D., Cetinic, I., Ferry, J., Fitzpatrick, E., Jones, B., Miller, P., Morton, S., Schaffner, R., Siegel, D., and Thunell, R. (2009). The toxin express: Rapid downward transport of the neurotoxin domoic acid in coastal waters. *Nature Geosciences*, 2(4):273–275.
- Shchepetkin, A. and McWilliams, J. (1998). Quasi-monotone advection schemes based on explicit locally adaptive dissipation. *Monthly Weather Review*, 126:1541–1580.
- Shchepetkin, A. F. and McWilliams, J. C. (2005). The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model. *Ocean Modelling*, 9:347–404.
- Shulman, I., Anderson, S., Rowley, C., DeRada, S., Doyle, J., and Ramp, S. (2010). Comparisons of upwelling and relaxation events in the Monterey Bay area. *Journal of Geophysical Research*, 115(C06016):12pp.
- Shulman, I., Kindle, J., Martin, P., deRada, S., Doyle, J., Penta, B., Anderson, S., Chavez, F., Paduan, J., and Ramp, S. (2007). Modeling of upwelling/relaxation events with the navy coastal ocean model. *Journal of Geophysical Research*, 112.
- Shulman, I., Rowley, C., Anderson, S., DeRada, S., Kindle, J., Martin, P., Doyle, J., Cummings, J., Ramp, S., Chavez, F., Fratantoni, D., and Davis, R. (2009). Impact of glider data assimilation on the Monterey Bay model. *Deep-Sea Research II*, 56:188 – 198.
- Singh, A., Krause, A., Guestrin, C., and Kaiser, W. (2009). Efficient informative sensing using multiple robots. *Journal of Artificial Intelligence Research*, 34:707–755.
- Smith, R. N., Chao, Y., Li, P. P., Caron, D. A., Jones, B. H., and Sukhatme, G. S. (2010a). Planning and implementing trajectories for autonomous underwater vehicles to track evolving ocean processes based on predictions from a regional ocean model. *International Journal of Robotics Research*, 29(12):1475–1497.
- Smith, R. N., Das, J., Heidarrson, H., Pereira, A., Caron, D. A., Jones, B. H., and Sukhatme, G. S. (2009). Implementation of an embedded sensor network for the coordination of slocum gliders for coastal monitoring and observation. In *WUWNet '09: Proceedings of the Fourth ACM International Workshop on UnderWater Networks*, pages 1–8, Berkeley, California. ACM.

- Smith, R. N., Das, J., Heidarsson, H., Pereira, A., Cetinić, I., Darjany, L., Ève Garneau, M., Howard, M. D., Oberg, C., Ragan, M., Schnetzer, A., Seubert, E., Smith, E. C., Stauffer, B. A., Toro-Farmer, G., Caron, D. A., Jones, B. H., and Sukhatme, G. S. (2010b). USC CINAPS builds bridges: Observing and monitoring the Southern California Bight. *IEEE Robotics and Automation Magazine, Special Issue on Marine Robotics Systems*, 17(1):20–30.
- Smith, R. N., Kelly, J., Chao, Y., Jones, B. H., and Sukhatme, G. S. (2010c). Towards improvement of autonomous glider navigation accuracy through the use of regional ocean models. In *Proceedings of the 29th International Conference on Offshore Mechanics and Arctic Engineering*, pages 597–606, Shanghai, China.
- Smith, R. N., Pereira, A., Chao, Y., Li, P. P., Caron, D. A., Jones, B. H., and Sukhatme, G. S. (2010d). Autonomous underwater vehicle trajectory design coupled with predictive ocean models: A case study. In *Proceedings of The IEEE International Conference on Robotics and Automation*, pages 4770–4777, Anchorage, Alaska.
- Smith, R. N., Schwager, M., Smith, S. L., Rus, D., and Sukhatme, G. S. (2011). Persistent ocean monitoring with underwater gliders: Towards accurate reconstruction of dynamic ocean processes. In *Proceedings of the International Conference on Robotics and Automation*, Shanghai, China.
- Smith, S. L. and Rus, D. (2010). Multi-robot monitoring in dynamic environments with guaranteed currency of observations. In *IEEE Conference on Decision and Control*, pages 514–521, Atlanta, GA.
- Smith, S. L., Schwager, M., and Rus, D. (2010e). Persistent robotic tasks: Monitoring and sweeping in changing environments. *IEEE Transactions on Robotics*. Submitted.
- Stommel, H. (1989). The slocum mission. *Oceanography*, 2:22–25.
- Taix, M., Souères, P., Frayssinet, H., and Cordesses, L. (2006). Path planning for complete coverage with agricultural machines. In Yuta, S., Asama, H., Prassler, E., Tsubouchi, T., and Thrun, S., editors, *Field and Service Robotics*, volume 24 of *Springer Tracts in Advanced Robotics*, pages 549–558. Springer Berlin / Heidelberg.
- Thomas, A. C., Brickley, P., and Weatherbee, R. (2009). Interannual variability in chlorophyll concentrations in the humboldt and california current systems. *Progress in Oceanography*, 83:386 – 392.
- Vu, Q. (2008). JPL OurOcean portal. <http://ourocean.jpl.nasa.gov/>. viewed September 2010.
- Warrick, J. A., DiGiacomo, P. M., Weisberg, S. B., Nezlin, N. P., Mengel, M. J., Jones, B. H., Ohlmann, J. C., Washburn, L., Terrill, E. J., and Farnsworth, K. L. (2007). River plume patterns and dynamics within the southern california bight. *Continental Shelf Research*, 27:2427?2448.
- Washburn, L., McClure, K. A., Jones, B. H., and Bay, S. M. (2003). Spatial scales and evolution of stormwater plumes in santa monica bay. *Marine Environmental Research*, 56(1 - 2):103–125.
- Webb Research Corporation (2008). <http://www.webbresearch.com/slocum.htm>. viewed March 2011.