# Trajectory Design for Autonomous Underwater Vehicles based on Ocean Model Predictions for Feature Tracking

Ryan N. Smith, Yi Chao, Burton H. Jones, David A. Caron, Peggy P. Li and Gaurav S. Sukhatme

**Abstract** Trajectory design for Autonomous Underwater Vehicles (AUVs) is of great importance to the oceanographic research community. Intelligent planning is required to maneuver a vehicle to high-valued locations for data collection. We consider the use of ocean model predictions to determine the locations to be visited by an AUV, which then provides near-real time, *in situ* measurements back to the model to increase the skill of future predictions. The motion planning problem of steering the vehicle between the computed waypoints is not considered here. Our focus is on the algorithm to determine relevant points of interest for a chosen oceanographic feature. This represents a first approach to an end to end autonomous prediction and tasking system for aquatic, mobile sensor networks. We design a sampling plan and present experimental results with AUV retasking in the Southern California Bight (SCB) off the coast of Los Angeles.

## 1 Introduction

More than three-fourths of our earth is covered by water, yet we have explored less than 5% of the aquatic environment. Autonomous Underwater Vehicles (AUVs) play a major role in the collection of oceanographic data. To make new discoveries and improve our overall understanding of the ocean, scientists must make use of these platforms by implementing effective monitoring and sampling techniques to study ocean upwelling, tidal mixing or other ocean processes. One emerging example of innovative and intelligent ocean sampling is the automatic and coordinated control of autonomous and Lagrangian sensor platforms [4].

As complex and understudied as the ocean may be, we are able to model and predict certain behaviors moderately well over short time periods. Expanding our modeling capabilities, and general knowledge of the ocean, will help us better exploit the resources that it has to offer. Consistently comparing model predictions with actual events, and adjusting for discrepencies, will increase the range of validity of existing ocean models, both temporally and spatially.

R.N. Smith and G.S. Sukhatme are with the Robotic Embedded Systems Laboratory, University of Southern California, Los Angeles, CA 90089 USA, e-mail: {ryannsmi,gaurav}@usc.edu
· Y. Chao and P.P. Li are with the Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109 USA, e-mail: {yi.chao,p.p.li}@jpl.nasa.gov
· B.H. Jones and D.A. Caron are with the Department of Biological Sciences, University of Southern California, Los Angeles, CA 90089 USA, e-mail: {bjones,dcaron}@usc.edu

The goal of this paper is to present an innovative ocean sampling method that utilizes model predictions and AUVs to collect interesting oceanographic data that can also increase model skill. Our motivation is to track and collect daily information about an ocean process or feature which has a lifespan on the order of a week. We use an ocean model to predict the behavior of an interesting artifact, *e.g.,* a fresh water plume, over a small time period, *e.g.,* one day. This prediction is then used as input to an algorithm that determines a sampling plan for the AUV(s). The AUV(s) are then retasked from a current mission or deployed. Afterward, the collected data is assimilated into the ocean model and an updated prediction is computed. A new sampling plan is created and the process repeats until the artifact is out of range or is no longer of interest.

We motivate the work from an oceanographic perspective and provide a realistic field application. Next, we briefly describe the ocean model and AUV used in this study. We discuss the waypoint selection algorithm and present results from a field implementation. We conclude with future research plans.

The work presented here serves as a proof of concept for the utilization of ocean model forecasts to design sampling missions for AUVs in particular, and aquatic mobile sensor platforms in general, to follow an ocean feature and collect data.

## 2 Oceanography Application and Ocean Model

Microscopic organisms are the base of the food chain: all aquatic life ultimately depends upon them for food. There are a few dozen species of phytoplankton and cyanobacteria that can create potent toxins when provided with the right conditions. Harmful algal blooms (HABs) can cause harm via toxin production, or by their accumulated biomass. Such blooms can cause severe illness and potential death to humans as well as to fish, birds and other mammals. The blooms generally occur near fresh water inlets, where large amounts of nutrient rich, fresh water is deposited into the ocean. This water provdes the excess food to support higher productivity and a *bloom* of microorganisms. It is of interest to predict when and where HABs may form, and which coastal areas they may affect. Harmful algal blooms are an active area of research along the western coast of the United States and are of large concern for coastal communities in southern California. The impact of HABs in this region can be seen in [7, 10]. With this motivation, we choose fresh water plumes as an ocean feature for which to design predictive tracking missions.

The predictive tool utilized in this study is the Regional Ocean Model System (ROMS) [9] - a split-explicit, free-surface, topography-following-coordinate oceanic model. We use ROMS because it is an open source ocean model that is widely accepted and supported throughout the oceanographic and modeling communities. Additionally, the model was developed to study ocean processes along the western U.S. coast which is our primary area of study.

Research is currently ongoing to update and improve ROMS for the Southern California Bight (SCB)[1] in an effort to characterize and understand the complex up-

[1] The SCB is the oceanic region contained within 32° N to 34.5° N and $-117°$ E to $-121°$ E

welling and current structure that exist and drive the local climate. The Jet Propulsion Laboratory (JPL) uses ROMS to provide nowcasts and hourly forecasts (up to 36 hours) for Monterey Bay, the SCB and Prince William Sound, see [6] for more information. The JPL version of ROMS assimilates HF radar surface current measurements, data from moorings, satellite data and any data available from AUVs operating in the area. Information regarding this specific version of ROMS and the data assimilation process can be found in [3].

## 3 Mobile Sensor Platform: AUV

The mobile sensor platform used in this study is a Webb Slocum autonomous underwater glider, as seen in Fig. 1. (http://www.webbresearch.com) The Slocum glider is a type of AUV designed for long-term ocean sampling and monitoring [8]. These gliders *fly* through the water by altering the position of their center of mass and changing their buoyancy. Due to this method of locomotion, gliders are not fast moving AUVs, and generally have operational velocities on the same order of magnitude as oceanic currents. The endurance and velocity characteristics of the glider make it a good candidate vehicle to track ocean features which have movements that are determined by currents, and that have a residence time on the order of weeks.

**Fig. 1** *He Ha Pe*, one of two USC Slocum gliders, flying a mission off the coast of Catalina Island.

We utilize autonomous gliders because our collaborative research group owns two of them, and hence field experiments can be readily performed. We have upgraded the communication capabilities of our vehicles to take advantage of our local wireless network; details on this can be found in the concurrent article, [5].

Extensive research has been done on glider dynamics and controller design, *e.g.,* see [2] and the references therein. Thus, we do not discuss these details nor the trajectory along which the glider travels. We assume here that the glider can successfully navigate from one location to another.

## 4 Trajectory Design

We now present an algorithm which generates the locations for the AUV to visit to follow the general movements of a fresh water plume through the ocean.

Considerable study has been reported on adaptive control of single gliders and coordinated multi-glider systems, see for example [4] and the included references. In these papers, the trajectories given to the gliders were fixed patterns (rounded polygons) that were predetermined by a human operator. The adaptive control component was implemented to keep the gliders in an optimal position, relative to the other gliders following the same trajectory. The difference between the method used in [4] and the approach described here, is that here the sampling trajectory is determined by use of the output of ROMS, and thus is, at first glance, a seemingly random and irregular sampling pattern. Such an approach is a benefit to the model and scientist alike. Scientists can identify sampling locations based upon ocean measurements they are interested in following, rather than setting a predetermined trajectory and hoping the feature enters the transect while the AUV is sampling. Model skill is increased by the continuous assimilation of the collected data; which by choice, is not a continuous measurement at the same location.

For a fresh water plume, the low salinity and density imply that this feature will propagate through the ocean driven primarily by surface currents. A plume may dissipate rapidly, but can stay cohesive and detectable for up to weeks; we assume the later case. It is of interest to track these plumes based on the discussion in Sect. 2 as well as in [1]. In addition to tracking the plume, it is also important to accurately predict where a plume will travel on a daily basis. The ROMS prediction capabilities for a plume are good, but model skill can significantly increase from assimilation of *in situ* measurements.

A single Slocum glider is not optimal for the task at hand, as it is built for endurance missions and traveling at low velocities. Hence, we can not expect it to be able to collect samples over the entire area of a potentially large plume. Thus, we restrict ourselves to visiting (obtaining samples at) at most two locations for each hour of sampling. The primary location that we are interesed in tracking is the centroid of the plume extent; analogous to the eye of the storm. Optimally, we would also like to gather a sample on the boundary of the plume. However, the glider may not be able to reach the plume centroid and a point on the boundary in a one hour time frame. With the given mission and the tools at hand, we present the following trajectory design algorithm.

### 4.1 Trajectory Design Algorithm Based on Ocean Model Predictions

We propose the following iterative algorithm for plume tracking utilizing ocen model predictions. This is the first known presentation of such a technology chain, and as such, is presented in a simplified manner. First, we assume the glider travels at a constant velocity $v$. Let $d$ be the distance in kilometers that the vehicle can travel in a given time. We neither consider vehicle dynamics nor the effect of ocean currents upon the vehicle in this study; these are areas of ongoing research. Also, we only consider a 2-D planar problem as far as the waypoint computation is concerned.

The input to the trajectory design algorithm is a set of points, $\mathscr{D}$ (referred to as drifters) that determine the initial extent of the plume, and hourly predictions of the location of each point in $\mathscr{D}$ for a set duration. For the points in $\mathscr{D}$, we compute the convex hull as the minimum bounding ellipsoid, $E_0$. The centroid of this ellipsoid, $C_0$, is the start point of the survey. Next, we consider the predicted locations of $\mathscr{D}$ after one hour, $\mathscr{D}_1$. The centroid of $\mathscr{D}_1$ is $C_1$; the centroid of the minimum bounding ellipsoid $E_1$. The algorithm computes $d_g(C_0,C_1)$, the geographic distance from $C_0$ to $C_1$. Given upper and lower bounds $d_u$ and $d_l$, resp., if $d_l < d_g(C_0,C_1) \leq d_u$, the trajectory is simply defined as the line $\overline{C_0C_1}$. If $d_g(C_0,C_1) \leq d_l$, the algorithm first checks to see if there exists a point $p \in E_1 \cup \mathscr{D}_1$ such that

$$d_l \leq d_g(C_0, p) + d_g(C_1, p) \leq d_u. \tag{1}$$

If such a point exists, the trajectory is defined as the line $\overline{C_0 p}$ followed by the line $\overline{pC_1}$. If the set of points $p \in E_1 \cup \mathscr{D}_1$ which satisfy Eq. 1 is empty, then the algorithm computes the locus of points, $\mathscr{L} = \{p^* \in \mathscr{L} | d_g(C_0, p) + d_g(p, C_1) = d\}$. This locus $\mathscr{L}$, by definition, defines an ellipse with focii $C_0$ and $C_1$. We then choose a random point $p^* \in \mathscr{L}$ as another location for sampling. Here, the trajectory is the line $\overline{C_0 p^*}$ followed by the line $\overline{p^* C_1}$. If $d_g(C_0, C_1) > d_u$, the algorithm aborts as the plume is traveling too fast for the chosen vehicle. The algorithm then repeats this process for the defined duration of tracking. This selection process of waypoints for the AUV to visit to track the plume is presented in Algorithm 1. The overall iterative process to

**Algorithm 1** Waypoint Selection Algorithm Based on Ocean Model Predictions

**Require:** Hourly forecasts, $\mathscr{D}_i$ for a set of points $\mathscr{D}$ defining the initial plume condition and its movement for a period of time, $T$.
  **for** $0 \leq i \leq T$ **do**
    Compute $C_i$, the centroid of the minimum bounding ellipsoid $E_i$ of the points $\mathscr{D}_i$.
  **end for**
  **while** $0 \leq i \leq T-1$ **do**
    **if** $d_l \leq d_g(C_i, C_{i+1}) \leq d_u$ **then**
      The trajectory is $\overline{C_i C_{i+1}}$.
    **else if** $d_g(C_i, C_{i+1}) \leq d_l$ and $\exists p \in E_i \cup \mathscr{D}_i$ such that $d_l \leq d_g(C_i, p) + d_g(p, C_{i+1}) \leq d_u$. **then**
      The trajectory is $\overline{C_i p}$ followed by $\overline{pC_{i+1}}$.
    **else if** $d_g(C_i, C_{i+1}) \leq d_l$ and $\{p \in E_i \cup \mathscr{D}_i | d_l \leq d_g(C_i, p) + d_g(p, C_{i+1}) \leq d_u\} = \emptyset$. **then**
      Compute $\mathscr{L} = \{p^* \in \mathscr{L} | d_g(C_0, p) + d_g(p, C_1) = d\}$, select a random $p^* \in \mathscr{L}$ and define the trajectory as $\overline{C_i p^*}$ followed by $\overline{p^* C_{i+1}}$
    **else if** $d_g(C_i, C_{i+1}) \geq d_u$ **then**
      Stop the algorithm. The plume is moving too fast for the selected AUV.
    **end if**
  **end while**

design an implementable plume tracking strategy based on ocean model predictions is given in Algorithm 2.

**Algorithm 2** Ocean Plume Tracking Algorithm Based on Ocean Model Predictions

**Require:** A significant fresh water plume is detected via direct observation or remotely sensed data such as satellite imagery.
  **repeat**
    A set of points ($\mathscr{D}$) is chosen which determine the current extent of the plume.
    Input $\mathscr{D}$ to ROMS.
    ROMS produces an hourly forecast for all points in $\mathscr{D}$.
    Input hourly forcast for $\mathscr{D}$ into the trajectory design algorithm.
    Execute the trajectory design algorithm (see Sect. 1).
    Uploaded computed waypoints to the AUV.
    AUV executes mission.
    The AUV sends collected data to ROMS for assimilation into the model.
  **until** Plume dissipates, travels out of range or is no longer of interest.

*Remark 1.* In the SCB, a vertical velocity profile of ocean current is generally not constant. Since the plume propagates on the ocean surface ($1-3$ m) and the glider operates at depths of $60-100$ m, it is not valid to assume that they are subjected to the same current regime, in both velocity and direction. Thus, it may be possible for a plume to *outrun* a slow-moving vehicle (*i.e.,* $d_g(C_i, C_{i+1}) \geq d_u$).

## 5 Implementation and Field Experiments in the SCB

The rainy season in southern California runs from November to March. During this time, storm events cause large runoff into local area rivers and streams, all of which empty into the Pacific Ocean. Two major rivers in the Los Angeles area, the Santa Ana and the Los Angeles River, input large fresh water plumes to the SCB. Such plumes have a high liklihood of producing HAB events. We deployed a Webb Slocum glider into the SCB on February 17, 2009 to conduct a month-long observation and sampling mission. For this deployment, the glider is programmed to execute a zig-zag pattern mission along the coastline, as depicted in Fig. 2, by navigating to each of the six waypoints depicted by the red and black bullseyes. Figure 2 also delineates the 20 m and 30 m isobaths, given by the green and red lines, respectively.

Unfortunately, weather and remote sensing devices did not cooperate to produce a rain event along with a detectable fresh water plume, so we were unable to retask the glider to track a real plume by use of Algorithm 2. Instead, we defined a pseudo-plume $\mathscr{D}$ with 15 initial drifter locations to demonstrate the proof of concept of this research. The pseudo plume is given by the blue line in Fig. 3.

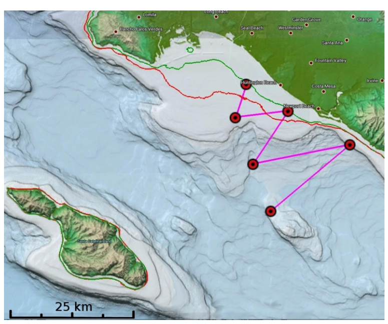


**Fig. 2** Preset waypoints, depicted with red and black bullseyes and the intended path of the glider given by the magenta line. Image created by use of Google Earth.

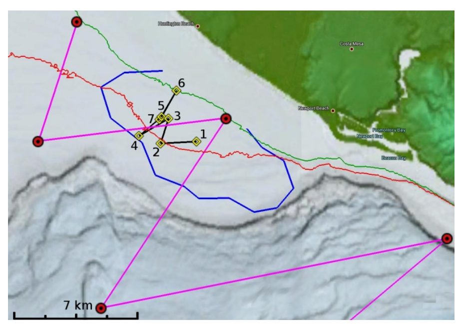


**Fig. 3** Plume (blue line), computed waypoints (yellow diamonds), and path connecting consecutive waypoints (black line). Image created by use of Google Earth.

The set $\mathscr{D}$ was sent to JPL and input to ROMS as the initial plume condition. The locations of the points in $\mathscr{D}$ were predicted for 15 hours. The initial time and location for the beginning of this retasking experiment coincided with predicted coordinates of a future glider communication. The pseudo-plume was chosen such that $C_0$ was near this predicted glider surfacing location.

Based on observed behavior for our vehicle during this deployment, we take $v = 0.75$ km/h, and initially defined $d_l = 0.5$ and $d_u = 0.8$. The hourly predictions were input to the trajectory design algorithm and a tracking strategy was generated. Due to slow projected surface currents in the area of study, the relative movement of the plume was quite small. To keep the glider from surfacing too often and to generate a more implementable trajectory, we opted to omit visiting consecutive centroids. Instead, we chose to begin at the initial centroid, then visit the predicted centroid of the plume after five, ten and 15 hours, $C_5$, $C_{10}$ and $C_{15}$, respectively. Between visiting these sites, the algorithm computed an additional waypoint for the glider to visit. These intermediate waypoints were chosen similarly to the $p^*$ defined earlier, with $d = 3.75$; the distance the glider should travel in five hours. This design strategy produced seven waypoints for the AUV to visit during the 15 hour mission. The waypoints are presented in Table 1.

Note that we include the initial centroid as a waypoint, since the glider may not surface exactly at the predicted location. Upon visiting all of the waypoints in Table 1, the glider was instructed to continue the sampling mission shown in Fig. 2. Figure 3 presents a broad overview of the waypoints in Table 1, along with a path connecting consecutive waypoints. The plume is delineated by the blue line and the waypoints are numbered and depicted by yellow diamonds. Note that the glider did not travel on the ocean surface during this experiment. Between waypoints, the glider submerges below a set depth and performs consecutive dives and ascents creating a sawtooth-shaped trajectory as its glide path.

**Table 1** Waypoints generated by the plume tracking algorithm. Waypoint numbers $1, 3, 5$ and $7$ are the predicted centroids of the pseudo-plume at hours $0, 5, 10$ and $15$, respectively.

| Number | Latitude (N) | Longitude (E) | Number | Latitude (N) | Longitude (E) |
|---|---|---|---|---|---|
| 1 | 33.6062 | -118.0137 | 5 | 33.6189 | -118.0349 |
| 2 | 33.6054 | -118.0356 | 6 | 33.6321 | -118.0257 |
| 3 | 33.6180 | -118.0306 | 7 | 33.6175 | -118.0361 |
| 4 | 33.6092 | -118.0487 | | | |

## 6 Results

In the study of path planning for field robots, planning the trajectory is usually less than half the battle, the real challenge comes in the implementation. This is exaggerated when dealing with underwater robots due to the complex environment. Next, we present results of an implementation of the designed sampling mission onto a Slocum glider operating in the SCB.

The waypoints given in Table 1 were computed under the assumption that the mission would be loaded onto the glider at a specific time and approximate geographic location. The glider arrived and communicated at the correct time and location, however, communication was aborted before the plume tracking mission could be uploaded. We were able to establish a connection two hours later at a different location, and successfully upload the mission file; this location is the red droplet labeled 1 in Fig. 4. We opted to not visit waypoint 1 based on the location of the glider and to get the glider back on schedule to track the plume. Figures 4 and 5 present magnified images of Fig. 3, where computed waypoints are the yellow diamonds and the red droplets are the actual locations visited by the glider.

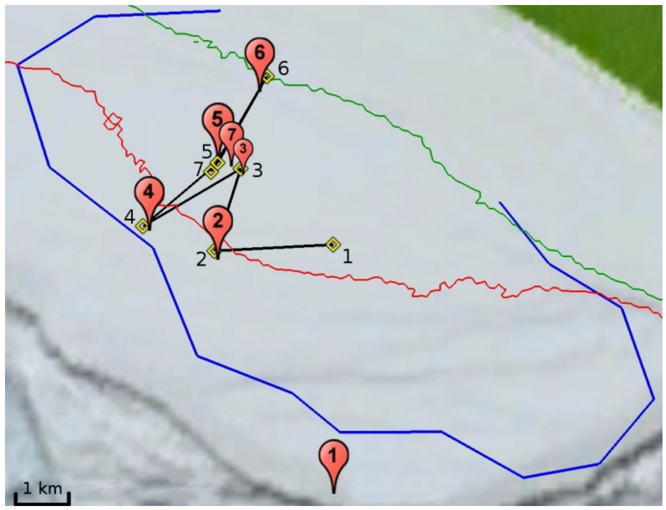


**Fig. 4** Computed waypoints (yellow diamonds) and actual glider locations (red droplets). Image created by use of Google Earth.

We were able to sucessfully generate a plan and retask a deployed glider to follow an ocean feature for 15 hours. It is clear from the data that consideration has to be made for glider dynamics and external forces from the ocean in the trajectory design algorithm. This is an area of active research. The motivation of this research is to follow plumes through the ocean via centroid tracking.

One element that we have neglected to discuss up to this point is that we have no metric for comparison. In particular, when we reach a predicted centroid, we do not have a method to check whether or not the plume centroid was actually at that location. We are planning experiments to deploy actual Lagrangian drifters to simulate a plume. This will give a comparison between the ROMS prediction and the actual movement of the drifters. This also provides a metric to determine the accuracy of the prediction and the precision of the AUV. Another component omitted from earlier discussion is time. When tracking a moving feature, a prediced waypoint contains time information as well as location. For this implementation, the glider began the mission at 0302Z and ended at 1835Z; a total time of 15.55 hours. Due to external influences, arrival at a few waypoints was not at the predicted time. Resolving this matter is contained within the addition of external forces, and is the subject of ongoing work.

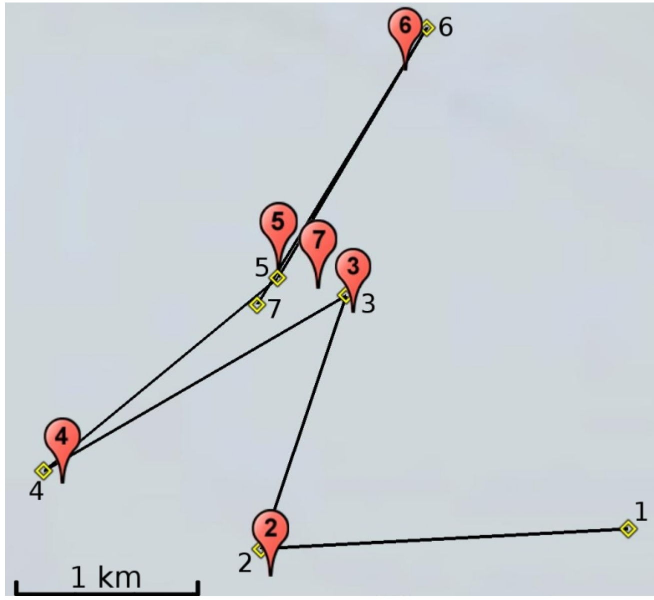


**Fig. 5** Computed waypoints (yellow diamonds) and actual glider locations (red droplets). Image created by use of Google Earth.

## 7 Conclusions and Future Work

Designing effective sampling strategies to study ocean phenomena is a challenging task and can be approached from many different angles. Here, we presented a method to exploit multiple facets of technology to achieve our goal. Utilizing an ocean model and an AUV, we were able to construct a technology chain which outputs a path to follow a fresh water plume centroid for a chosen period of time. The successful field experiment presented here required the cooperation and communication between many individuals. Retasking an autonomous glider remotely while it is in the field involves patience, determination and many resources. In a period of less than two hours, we were able to decide to retask the glider, delineate a plume in the ocean, use ROMS to generate a prediction, generate an implementable tracking strategy, create a glider mission file and have it ready to upload to the glider. This paper has demonstrated that we have implemented the collaboration and technology chain required to perform complex field experiments. The work now is to improve upon the waypoint generation algorithm and extend it to design implementable 3-D trajectories.

The main implementation issue is the ability of the glider to accurately navigate to a given waypoint. This is a direct result of the waypoint selection algorithm only solving the 2-D problem, and ignoring the dynamics of the glider and the complex ocean environment. Details on how to implement robustness and generate more complex sampling missions are outside the scope of this paper. Areas of ongoing research include plans to incorporate the kinematic and dynamic models of the glider and extend this from a planar to a 3-D motion planning algorithm. Also, we plan to incorporate a 3-D current output of ROMS to plan a trajectory that exploits the currents to aid the locomotion of the glider. A more immediate step is to incorporate multiple AUVs, which leads to the development of an optimization criterion on which vehicle is best suited for a certain mission or to visit a chosen waypoint. A long-term goal is to facilitate autonomy for the entire system, leaving the human in the control loop as a fail-safe.

**Acknowledgements** This work was supported in part by the NOAA MERHAB program under grant NA05NOS4781228, by NSF as part of the Center for Embedded Networked Sensing (CENS) under grant CCR-0120778, by NSF grants CNS-0520305 and CNS-0540420, by the ONR MURI program under grant N00014-08-1-0693, and a gift from the Okawa Foundation. The ROMS ocean modeling research described in this publication was carried out by the Jet Propulsion Laboratory (JPL), California Institute of Technology, under a contract with the National Aeronautics and Space Administration (NASA).

The authors acknowledge Carl Oberg for his work with glider hardware making field implementations possible and simple. We thank Ivona Cetinic for her efforts in the glider deployment and communications during the missions.

## References

1. I. Cetinic, B.H. Jones, M.A. Moline, and O. Schofield. Resolving urban plumes using autonomous gliders in the coastal ocean. *Journal of Geophysical Research/Americal Geophysical Union*, 2008.
2. J.G. Graver. *Underwater Gliders: Dynamics, Control and Design*. PhD thesis, Princeton University, 2005.
3. Zhijin Li, Yi Chao, James C. McWilliams, and Kayo Ide. A three-dimensional variational data assimilation scheme for the regional ocean modeling system. *Journal of Atmospheric and Oceanic Technology*, 25:2074–2090, 2008.
4. D. Paley, F. Zhang, and N.E. Leonard. Cooperative control for ocean sampling: The glider coordinated control system. *IEEE Transactions on Control Systems Technology*, 16(4):735–744, 2008.
5. A. Pereira, H. Heidarsson, D.A. Caron, B.H. Jones, and G.S. Sukhatme. An implementation of a communication framework for the cost-effective operation of slocum gliders in coastal regions. In *Proceedings of The* $7^{th}$ *International Conference on Field and Service Robotics*, Cambridge, MA, July, 2009. Submitted.
6. Q. Vu. JPL OurOcean portal. http://ourocean.jpl.nasa.gov/, 2008. viewed February 2009.
7. A. Schnetzer, P. E. Miller, R. A. Schaffner, B. Stauffer, B. Jones, S. B. Weisberg, P. M. DiGiacomo, W. M. Berelson, and D. A. Caron. Blooms of pseudo-nitzschia and domoic acid in the san pedro channela and los angeles harbor areas of the southern california bight. *Harmful Algae/Elsevier*, 6(3):372–387, 2003-2004.
8. O. Schofield, J. Kohut, D. Aragon, E.L. Creed, J.G. Graver, C. Haldman, J. Kerfoot, H. Roarty, C. Jones, D.C. Webb, and S. Glenn. Slocum gliders: Robust and ready. *Journal of Field Robotics*, 24(6):473–485, 2007.
9. A.F. Shchepetkin and J.C. McWilliams. The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model. *Ocean Modeling*, 9:347–404, 2005. doi:10.1016/j.ocemod.2004.08.002.
10. E.S. Wood, A. Schnetzer, C.R. Benitez-Nelson, C. Anderson, W. Berelson, J. Burns, D.A. Caron, J. Ferry, E. Fitzpatrick, B. Jones, P.E. Miller, S.L. Morton, and R. Thunell. The toxin express: Rapid downward transport of the neurotoxin domoic acid in coastal waters. *Nature Geosciences*, 2009.