

## 2.10 Multiscale Actuated Sensing (MAS)

Multiscale actuated sensing research activities have focused on the core algorithmic challenges of incorporating rich image media and mobility as key components of environmental sensing systems, in particular: planning, design and calibration, and vision-based sensing, spatial sampling, and tracking.



Figure 8. Deployed Glider resting at the surface.

### Planning, Design and Calibration

Trajectory design for Autonomous Underwater Vehicles (AUVs; Figure 8) is of great importance to the oceanographic research community. We consider the use of ocean model predictions to determine the locations to be visited by mobile sensor platforms. The platforms, in turn, provide near-real time, in situ measurements back to the model to increase the skill of future predictions. We have expanded upon our previous proof-of-concept results and demonstrated open-ocean implementation with multiple vehicles on data-gathering campaigns in the Southern California Bight (SCB). We have also investigated the impact of the use of predicted ocean currents on the trajectory design process for underwater vehicles. We take an Informative Path Planning approach to the robot path-planning problem. We aim to choose measurements that best describe a scalar field of interest (e.g. temperature). Given a model for the underlying scalar field we are able to compute the covariance and entropy

of the field given a hypothetical set of sample locations. Using this model of the underlying field, we maximize the mutual information of the un-sensed locations with respect to the sensed locations, resulting in good AUV paths. A coastal communication system provides a fast, safe, affordable way of communicating with the underwater vehicles through a RF-modem base-station network installed in the Southern California region. Using this method we measure a 6x improvement (24x with compression) in file transfer speed, compared to Iridium. We determined experimentally that the feasible communication range using this method is approximately 12 km. Experiments with a month-long deployment in the SCB provide evidence that sea state is a factor that affects range due to the low antenna position. In collaboration with MBARI, we have used data available from a range of sources, including ocean models, remote sensing satellites, moorings, and on-shore instruments to predict the trajectory of a patch of water. In experiments targeting the Monterey Bay (a biologically diverse bodies of water that experiences extreme "red-tide" blooms), an ideal location for algal bloom studies. An advection analysis of bloom "hotspots" from October 2007 and October 2008 data in the Monterey Bay shows an example of how such predictions can be used to plan survey missions for AUVs. Finally, our recent work has demonstrated that visual and inertial sensors, in combination, can provide very accurate estimates of the ego-motion of a robotic sensing platform. The accuracy of the motion estimate depends, however, on proper calibration of the transform between the camera and the inertial measurement unit (IMU). Un-modeled calibration errors will introduce biases in the estimation process, degrading overall localization performance —sometimes dramatically. Although accurate calibration is critical, many existing camera-IMU calibration techniques are difficult, time-consuming and require additional complex apparatus. We have developed algorithms to circumvent many of these problems. Using these techniques we are able to accurately predict drift in the (purely inertial) estimates of AUV position.

### Vision-based Sensing, Spatial Sampling, and Tracking

There exist many biological sensing applications where direct measurement is impossible, invasive, or time consuming. For example, measuring the presence/absence of birds at a feeder station currently requires a human to watch a camera pointed at the feeder, identifying when birds arrive and leave. Similarly, measuring CO<sub>2</sub> flux from a plant requires placing the plant inside a growth chamber, destructively modifying the environment. We use imagers as biological sensors by constructing a procedure that uses images to obtain approximate measurements of such phenomena. This procedure, composed of state-of-the-art computer vision, image processing, and statistical learning algorithms, is evaluated in the context of two specific applications: automatically detecting and segmenting out animals from a natural scene, and pollinator counting. Our approach to bird localization is to first localize by exploiting the fixed view deployment scenario, and cluster detections based on the similarity of appearance and spatiotemporal coherence. We introduce a discrepancy measure that has several useful properties for multi-view categorization in the context of surveillance and monitoring applications. For pollinator counting, the procedure roughly consists of three independent pieces. For each frame, we first localize the region of interest as specified by the user. This effectively discards a large fraction of each frame that is uninteresting, by definition, and would only serve to confuse further processing. Next, we detect and localize the target as it occludes the region of interest. Finally, using the

match value and location of a potential match produced by the target detection, we track the target's motion over time, leveraging the inherent temporal correlation between frames to discard false detections. We also propose a method for classifying the vegetation types in an aerial color infrared image. Different vegetation types do not only differ in color, but also in texture. We study the use of four Haralick features (energy, contrast, entropy, and homogeneity) for texture analysis, and then perform the classification using the One-Against-All (OAA) multiclass Support Vector Machine (SVM). Whether using imagers or other modalities, detecting, localizing, and tracking targets are different aspects of monitoring that have wide applications. We propose the use of a new random set theory (RST) approach to these problems. RST is a theory developed on random sets, which generalizes probability theory to the set domain. We use this theory to model sensor failures, lost connections, noise, and clutter. In sensor networks, for dynamic events, classical joint estimation/detection/tracking multi-sensor and multi-target algorithms are often hybrids of both analytical and ad-hoc approaches at various levels. The intricacies of the resulting solutions when the number of targets and sensors may vary randomly often obscure design intuition and leave many design choices to a largely trial and error based approach. By treating multi-target and multi-sensor cases jointly, RST is able to provide a systematic framework for rigorous mathematical analysis. A rigorous statistical framework has been developed for RST that includes concepts as: ML, Bayesian filtering, data fusion, and Cramer-Rao Bound.

# MAS 01 Autonomous Underwater Vehicle Path Planning and Navigation Based upon Ocean Model Predictions

## MAS 01.1 Overview

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 selected locations to collect data with scientific merit. This planning is not only determining how to get from one point to another, but also in determining which locations should be visited to best understand the feature under investigation. In this study, we examine the use of ocean model predictions to determine the locations to be visited by an AUV, and aid in planning the trajectory that the vehicle executes during the sampling mission. The objectives are: a) to provide near-real time, in situ measurements to a large-scale ocean model to increase the skill of future predictions, b) to utilize ocean model predictions as a component in an end-to-end autonomous prediction and tasking system for aquatic, mobile sensor networks, and c) to investigate the impact of utilizing ocean model predictions in the trajectory design for autonomous underwater gliders to improve their navigational accuracy.

## MAS 01.2 Approach

As an overview, the basic mission plan studied in this paper is to track and collect daily information about dynamically evolving ocean processes or features. First, we identify a feature of interest in the Southern California Bight (SCB) via direct observation or by use of remotely sensed data. The SCB is the oceanic region contained within 32° N to 34.5° N and -117° E to -121° E. We then use an ocean model to predict the behavior of this feature, e.g., HAB, over a short time period, e.g., one day. This prediction is used to generate a sampling plan for deployed gliders that steers the vehicles to regions of scientific interest based upon the given feature. Throughout execution of the sampling plan, the collected data are transmitted and assimilated into the ocean model. A new prediction is generated and the process is repeated until the feature dissipates, or is no longer of interest. A long term goal is to collectively implement this and other path planning algorithms, e.g., (Das et al. 2010) and (Binney et al. 2010), via an embedded sensor network (see (Smith et al. 2010) and (Pereira et al. 2009)) to enable real-time, optimal path planning and trajectory design, based on ocean model predictions and in situ measurements gathered by a fleet of mobile sensor platforms to further our understanding of interesting and evolving ocean features and phenomena.

The motivation for using predictive capabilities to design trajectories with the intent of tracking an evolving ocean feature is derived from a practical problem that exists in many coastal communities around the world, and in particular, Southern California. As the rate of urbanization in coastal communities continues to increase, land use and land cover (e.g., a significant increase in impervious surfaces) in these areas are permanently altered. This alteration affects both the quantity of freshwater runoff, and its particulate and solute loadings, which has an unknown impact (physically, biogeochemically, biologically and ecologically) on the coastal ocean. One documented result of these impacts is an increase in the occurrence of algal and phytoplankton blooms. Such biological phenomena are a primary, collaborative research interest of the authors. In particular, we are interested in the assessment, evolution and potential prediction of the blooms that have the potential to include harmful algal species (i.e., Harmful Algal Blooms (HABs)).

The processes leading to the onset, evolution and dissemination of HAB events are still under investigation. The recent development and implementation of an embedded sensor network along the Southern California coast (see e.g., (Smith et al. 2010), with further details in (Pereira et al. 2009)) provides necessary sensors and infrastructure to facilitate an in-depth, multi-faceted investigation of physical, chemical and biological processes related to HABs in addition to the impacts resulting from urbanization and climate change. Pivotal components of this network are the mobile sensor platforms in the form of autonomous Slocum gliders. Based on their deployment longevity, coupled with the use of multiple gliders, these vehicles can provide an extended spatio-temporal series of observations. Additionally, gliders can be utilized to monitor and track dynamically evolving ocean features that have a lifespan on the order of weeks, i.e., HABs and freshwater runoff plumes.

The work conducted during this reporting period expands upon the previous proof-of-concept results and demonstrates open-ocean implementation. Further, we investigate the impact of the use of predicted ocean currents on the trajectory design process for the vehicle.

## MAS 01.3 System(s) Description and/or Experiments

The predictive tool utilized in this study is the Regional Ocean Model System (ROMS) - a split-explicit, free-surface, topography-following-coordinate oceanic model. ROMS 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. The model solves the primitive equations using the Boussinesq and hydrostatic approximations in vertical sigma (i.e., topography following) and

horizontal orthogonal curvilinear coordinates. ROMS uses innovative algorithms for advection, mixing, pressure gradient, vertical-mode coupling, time stepping, and parallel efficiency.

The version of ROMS used in this study is compiled and run by the Jet Propulsion Laboratory (JPL), California Institute of Technology. The JPL provides ROMS hindcasts, nowcasts and hourly forecasts (up to 36 hours) for the SCB. The JPL version of ROMS assimilates HF radar surface current measurements, data from moorings, satellite data and any data available from sensor platforms located or operating within the model boundary. This ROMS utilizes a nested configuration, with increasing resolution covering the U.S. west coastal ocean at 15 km, the southern California coastal ocean at 5 km, and the SCB at 1 km. The three nested ROMS domains are coupled online and run simultaneously exchanging boundary conditions at every time step of the coarser resolution domain. In addition to the 1 km output, a resampled 2.2 km resolution output is produced, which correlates to the assimilated HF radar grid resolution. This study utilizes this 2.2 km output for our computations. Current velocity predictions are provided at depths of 0, 5, 10, 15,...,2000 m, as the bathymetry permits. A concern with any model is the accuracy and precision of the predictions. Specifically, we are concerned with the spatial structure of the predicted current velocities. ROMS primarily assimilates surface velocities from HF radar data, and it is assumed that forecasting for near-surface velocities are reasonable. However, in areas containing a large vertical complexity of currents, particularly in shelf break regions, ROMS has demonstrated poor prediction capabilities. It is an area of active research to reduce the uncertainty, improve the performance, and improve the quality and utility of ROMS forecasts. The interaction between this research and ROMS improvement is a two-way street. We need the predictions to design efficient and effective trajectories, and ROMS utilizes the feedback from field deployments to assess the validity of each prediction.




Figure 1. Iterative algorithm output to determine the bearing and range at which the vehicle should dead reckon. The blue line is the initial path if no ocean current existed. The black line is the first iteration attempt. The point labeled Aim is the dead reckoning waypoint at which the vehicle heads to end up at the point Final, following the green line.

The mobile sensor platform used in this study is a Webb Slocum autonomous underwater glider, as seen in Fig. 1. The Slocum glider is a type of AUV designed for long-term ocean sampling and monitoring. 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 (~ 1 km/hr). 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 have the capability to remain of scientific interest for weeks at a time. An example mission for a Slocum glider is a maximum depth along with a set of preprogrammed waypoints ( $W_1, \dots, W_n$ ) that define the mission. As previously mentioned, a

typical trajectory is a sawtooth-shaped path. Each down/up cycle is referred to as a yo, and we define a segment to be the composition of multiple yos that begins with a dive from the surface and ends with a surfacing. Each time at the surface, the glider acquires a GPS location. The present location of the vehicle ( $L$ ) 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  $L$  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 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. In general, for large-scale, open-ocean, sampling and monitoring missions, as the gliders were designed for, this type of mission planning and execution is reasonable. Specifically, accuracy and precision are not a top priority, and open-ocean currents are relatively uniform both spatially and temporally. Thus, an estimation of the

local currents based on the error observed in the execution of the prior segment is reasonable. However, in a coastal region, with the intent to track and monitor an evolving ocean feature, accuracy becomes increasingly important, and current structures in this regime vary significantly both spatially and temporally. It is with these motivations that we investigate the use of 4-D ocean model velocity predictions to aid in the trajectory design process for gliders.

Additionally, we investigate the use of ocean model predictions of ocean currents to incorporate into the trajectory design process to increase navigational accuracy of the vehicles. To do this we first computed the dead-reckoning error for our vehicles based on the on-board sensor suite. We fused the measurements from the simulated attitude sensor and pressure sensor in an unscented Kalman filter (UKF) to estimate the position, attitude and velocity of the vehicle over time. The UKF is a Bayesian filtering algorithm which employs a statistical local linearization procedure to propagate and update the system state. For our simulations, we assume that the glider follows a nominal linear sawtooth trajectory, and that the vehicle angular rotation rate and linear acceleration are driven by white, zero-mean Gaussian noise processes. Once we computed the error inherent in the vehicle and its sensors, as if it were navigating in a motionless fluid, we were able to analyze the effect ocean currents on vehicle navigation. Field trials were performed to demonstrate the advantage of path planning for autonomous glides using predictive ocean models, such as ROMS.

#### MAS 01.4 Accomplishments

Given a freshwater plume, we are interested in designing missions for gliders to track and sample along the path of the centroid and boundary of the feature. We assume that we have at least two vehicles to perform the missions; one centroid tracker and one boundary tracker. For safety concerns, we restrict a glider to surface no more than once every four hours. Since the basic idea is to track the feature for many days, while assimilating collected data into the model, and the accuracy of the model prediction degrades with time, we choose to plan a 16 hour mission for each day, with  $T \in \{0, \dots, 16\}$  denoting the time within a given mission.

The starting location  $L$  of each vehicle is known, and we assume that the prediction of the plume evolution is accurate. There is no adaptive behavior incorporated during the execution of the planned path. The idea is to improve the collection of scientific data by predicting the best locations to send a glider, while also providing feedback to JPL on the accuracy of ROMS. In the long run, both communities will benefit. Additionally, we assume that the glider travels at a constant horizontal speed  $v$  km/h;  $dh$  km is the distance traveled in  $h$  hours.

The input to the centroid-tracking, waypoint-generation algorithm is a set of points,  $D$  (referred to as drifters) that determine the initial extent of the plume  $D_0$ , and hourly predictions ( $D_i, i \in T$ ) of the location of each point in  $D$ . For the points in  $D_i$ , we compute the convex hull as the minimum bounding ellipsoid,  $E_i$  for  $i \in T$ . We consider the predicted locations of  $D_0$  after four hours,  $D_4$ . The centroid of  $D_4$  is  $C_4$ ; the centroid of  $E_4$ . The algorithm computes  $d_g(L, C_4)$ , the geographic distance from  $L$  to  $C_4$ . Given upper and lower bounds  $d_u$  and  $d_l$ , resp., if  $d_l < d_g(L, C_4) \leq d_u$ , the generated waypoint is  $C_4$ , and the path is simply defined as the line from  $L$  to  $C_4$ . If  $d_g(L, C_4) \leq d_l$ , the algorithm first checks to see if there exists a point  $p \in E_4 \cap D_4$  such that  $d_l \leq d_g(L, p) + d_g(C_4, p) \leq d_u$ . If such a point exists, the algorithm generates two waypoints ( $p$  and  $C_4$ ) and the path is defined as the line from  $L$  to  $p$  followed by the line from  $p$  to  $C_4$ . If  $\{p \in E_4 \cap D_4 \mid d_l \leq d_g(L, p) + d_g(C_4, p) \leq d_u\} = \emptyset$ , then the algorithm computes the locus of points,  $L = \{p \in L \mid d_g(L, p) + d_g(p, C_4) = d_4\}$  and selects a point at random  $p^* \in L$  as another waypoint. Here, the path is the line from  $L$  to  $p^*$  followed by the line from  $p^*$  to  $C_4$ . If  $d_g(L, C_4) > d_u$ , the algorithm generates a waypoint  $C_w$  in the direction of  $C_6$ , such that  $d_g(L, C_w) = d_4$ . The location of the vehicle  $L$  is updated to  $C_4$  or  $C_w$  and the process is iterated until  $T \leq 16$ .

Similarly, we define the boundary-tracking, waypoint-generation algorithm. We begin with the same predictions as above, and define  $P_i$  to be the polygon formed by connecting the points  $D_i$  for  $i \in T$ . We define  $B(p, r)$  to be the 2-dimensional disc of radius  $r$ , about  $p$ . This algorithm first computes  $N = B(L, d_4) \cap P_4$ . If  $N \neq \emptyset$ , the generated waypoint  $B_4$  is a random selection of one of the intersection points. If  $N = \emptyset$ , the generated waypoint  $B_4$  is that precise intersection point. If  $N = \emptyset$ ,  $B_4$  is computed such that  $d_g(L, B_4) = d_4$  and  $AZ(L, B_4)$  is the average azimuth of  $D_i$  for the considered four hour time period. We reassign  $L = B_4$ , and the algorithm is repeated.

These two aforementioned algorithms generate the waypoints defining a mission for a vehicle to track an evolving ocean feature. The first version of the centroid tracking algorithm was presented in (Smith et al. 2009a). A preliminary, multi-vehicle simulation result from the boundary tracking algorithm is given in (Smith et al. 2009b). These initial versions of the algorithms were designed to solve the planar path planning problem. Since a glider does not travel on the ocean surface during deployments, and the vertical distribution of velocity cannot be assumed constant, we extended these prior efforts to consider the 3-D path planning scenario including external forcing in (Smith et al. 2010c).

The Ocean Plume Tracking Algorithm Built On Ocean Model Predictions (OPTA-BLOOM-Pred) takes as input  $L$ , the output waypoints of the boundary and centroid tracking algorithms (i.e.,  $C_i$  and  $B_i$  for  $i = 4, 8, 12, 16$ , resp.), as well as the 4-D (three spatial and one temporal) ROMS velocity predictions. The output of the algorithm is an alternate

waypoint e.g.,  $C_{4alt}$  for each input, e.g.,  $C_4$ . This alternate waypoint is the location to which the vehicle should dead reckon, so that it arrives at  $C_4$ , given ROMS predictions of the velocity field during the time the vehicle maneuvers from  $L$  to  $C_4$ . Figure 1 displays example start, end and dead reckon waypoints for one segment of a mission given by the points labeled Start, Final and Aim, respectively.

At this point in the study, we do not consider the complete dynamic equations of motion of the vehicle, but initially assume a simple kinematic model, where the total inertial velocity of the vehicle is the sum of the vehicle's body-fixed velocity (relative to motionless fluid) and the inertial velocity of the fluid itself (ocean currents). As an initial assumption, the glider trajectory is parameterized by a cosine curve, containing an integer number of periods (i.e., starts and ends at the surface), that is estimated from data collected over the course of multiple field deployments. Note that OPTA-BLOOM-Pred can accept any time-discretized, periodic vehicle trajectory. The glider's velocity along this parameterized trajectory is estimated from a glider simulator provided by the vehicle manufacturer, then normalized so that the average horizontal speed is  $v$  km/h. Based upon favorable proof-of-concept deployment results utilizing these algorithms, work is ongoing to incorporate the dynamic equations of motion to produce vehicle trajectories.

Given the inputs above and starting locations  $L_C$  and  $L_B$  for the gliders, OPTA-BLOOM-Pred computes a sinusoidal trajectory from  $L_C$  to  $C_4$ . When projected to the sea surface, this initial path minimizes the Euclidean distance between  $L_C$  to  $C_4$ . Based on the assumed speed of the glider, an estimated completion time  $T$  is computed. This trajectory,  $P(x,y,z,t)$ , is discretized into 30 second time steps  $T$  and the predicted current is superimposed iteratively at each step.

Tracking a dynamically evolving feature is highly time-dependent. Reaching a specified location has an associated, or accepted time of arrival. For a glider, time is lost at the surface due to GPS localization and communication. Thus, in addition to safety, it is in our best interest to minimize surfacings. To this end, in the computation of the centroid tracking waypoints, we additionally consider predicted feature centroids at \$5, 6, 7\$ and \$8\$ hours from  $L_C$ . The same iterations are applied, and a time to goal is estimated for  $C_{Jalt}$ ,  $J = 5, 6, 7, 8$ . The waypoint with the path that has an estimated end time closest to the time the centroid will be in that predicted location is chosen as the waypoint at which to surface. Note that this procedure may reduce the total number of surfacings over the duration of the mission.

For the boundary tracking mission, we implement the same iterations to compute  $C_{Jalt}$ ,  $J = 4, 8, 12, 16$ . However, we do not consider the optimization described in the previous paragraph. As the boundary is not a single point location, it is not as critical to match its precise spatial and temporal location. Additionally, we expect the prediction of the boundary of the feature to be less accurate since it is predicted by the evolution of an unconnected, discrete number of points. This process produces 4 boundary waypoints, and at most 4 centroid waypoints. The overall OPTA-BLOOM-Pred is given below.

Ocean Plume Tracking Algorithm Built On Ocean Model Predictions (OPTA-BLOOM-Pred)

REQUIRE: A significant fresh water plume is detected via direct observation or remotely sensed data.

REPEAT:

A set of points (D) is chosen which determine the current extent of the plume.

Input D to ROMS.

ROMS produces an hourly forecast for all points in D.

Input hourly forecast for D into centroid and boundary tracking algorithms.

Execute waypoint generation, centroid and boundary tracking algorithms.

Compute alternate waypoints that coordinate spatial and temporal movement of the feature.

Compute the alternate waypoints at which the vehicle aims, to arrive at the prescribed goal location.

Uploaded computed alternate 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.

We present a simulation result for the implementation of OPTA-BLOOM-Pred. In this example, a proxy feature of interest is delineated off the coast of Newport Beach, CA to emulate a fresh water plume.

In Fig. 2, we present the results of a 16 hour feature tracking mission. Figure 2A shows an overview of the area where the proxy feature was delineated. Figures 2B-F present the proposed mission in 4-hour time increments. The starting location for both gliders is given by the black star. Boundary tracking waypoints are given by red triangles, centroid tracking waypoints are given by yellow diamonds, the centroid of the current state of the plume is given by the orange dot, and the green dots represent the alternate waypoint at which the gliders dead reckon. In Fig. 2E, the path tracking the centroid remains the same as that presented in Fig. 2D. Here, OPTA-BLOOM-Pred computed that it was better to surface after eight hours rather than after four hours. Although the glider did not start directly at the centroid, by hour eight, it is supposed to surface directly at its predicted location. This occurs again at  $T = 16$ . We remark that in Figs. 2D-F, it appears that the computed centroid does not appear to be in a centralized location. This

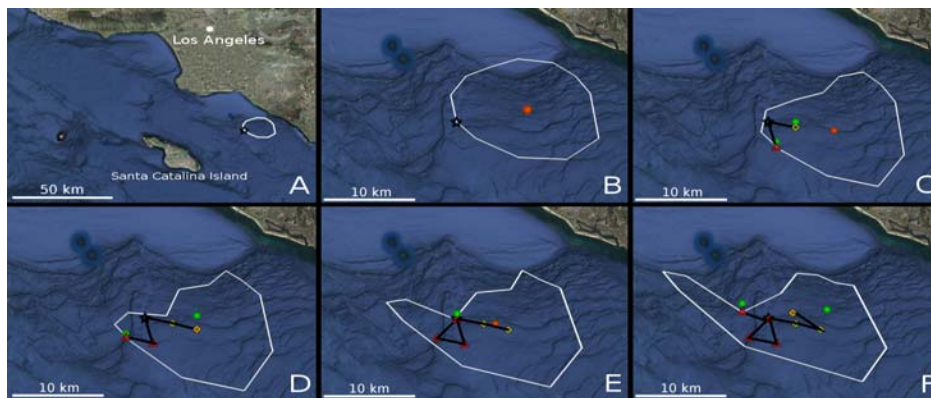


Figure 2. OPTA-BLOOM-Pred simulation results. Tracking of a proxy feature for 16 hours off the coast of Newport Beach, Los Angeles, CA. Images created by use of Google Earth.

is a result of computing the centroid based on the convex hull of the set of drifters. Since the boundary of this feature evolves into a non-convex shape, the centroid migrates toward the concave boundary.

The boundary tracking trajectory actually doubles back on itself in hour 12 as the predicted movement of the feature puts the boundary in a similar location at both hour zero and hour 12.

It is interesting to note that neither of these paths would be defined by a human operator, and at first glance, are seemingly somewhat random. However, they are designed to track an evolving feature that does not move following set patterns. This tracking appears to work well and gives promise to the implementation of trajectories generated by model predictions. In each of Figs. 2B-F, we can see that the difference between dead reckoning location and goal location can be vastly different based on location and time. The alternate waypoints at which to dead reckon heavily depend on the accuracy of ROMS. The field trials presented in the following section provide an initial investigation into the accuracy of ROMS and the validity of this method of technology fusion.

Based on previous deployment experience with two Slocum gliders, we estimated a median error of 1.1 km between the actual surfacing location and the goal waypoint over an average trajectory length of 2 km. For missions such as feature tracking, as described in (Smith et al. 2009a) and (Smith et al. 2010b), we would like to have better navigational accuracy. By fusing the measurements of simulated attitude and pressure sensors in an unscented Kalman filter, we estimated the position, attitude and velocity of the vehicle over time. Based on the specifications of the instruments on-board the glider, we were also able to propagate the uncertainty of the vehicle's position, attitude and velocity. In this study, we were primarily interested in the uncertainty of the final surfacing location after execution of a 2 km trajectory. From the UKF, we estimated that the three-sigma dead reckoning error for our trajectory is an ellipsoid that has a semi-major axis of 582.6 m, a semi-minor axis of 41.1 m, and an area of about 75,000 m<sup>2</sup>. As seen in Figs. 3 and 4, this corresponds to a cross-track error of approximately 600 m, at an angle of 16.7°, in the worst case. Thus, we can reasonably expect that the best achievable accuracy for our Slocum gliders, which dead reckon by use of only an attitude and a pressure sensor, is on the order of 0.5 km for every 2 km traversed. We reiterate that this is roughly half the error that we have seen throughout multiple deployments. However, this estimation matches well with the experimental errors for the deployments that used ROMS predictions for the in situ ocean current rather than the depth-averaged current measurements computed on-board the glider published in (Smith et al. 2010c); 500 m over a 1.7 km trajectory. Although this may be coincidental, the data provide motivation

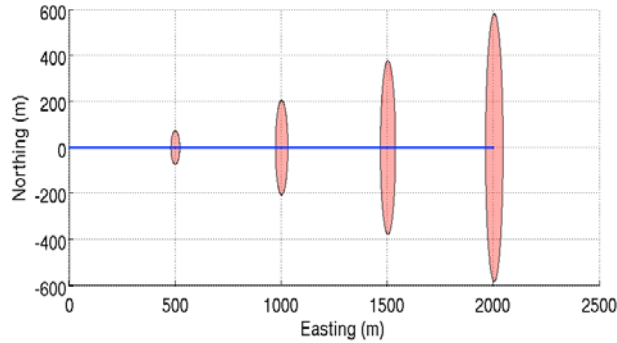


Figure 3. Downward looking view of a 2 km glider trajectory (blue line) with 3-sigma uncertainties at 500, 1000, 1500 and 2000 meters.




Figure 4. Three-dimensional view of a 2 km glider trajectory (blue line) with 3-sigma (red ellipsoids) uncertainties at approximately 500, 1000, 1500 and 2000 meters.

for further investigation into the incorporation of ROMS predictions into the trajectory design and basic navigation for autonomous gliders. This is an area of active research being carried out in collaboration with the JPL.

### MAS 01.5 Future Directions

This research is a small portion of a large effort that is preparing a region-wide, coastal-ocean survey called Bight 2010. This is a multi-facility, comprehensive study of the SCB. Bight 2010 is planned for February 2010 through April 2010, and we are planning to implement our algorithms on at least two deployed gliders to track and monitor an actual fresh water plume or algal bloom. During all subsequent deployments, we will utilize ROMS predictions in the trajectory design to improve the navigational accuracy of the vehicles. Comparative studies will be conducted where one glider operates normally while the other incorporates ROMS predictions while executing the same mission. These deployments will provide data necessary to assess the validity and accuracy of ROMS predictions in select areas throughout the SCB. These data can be used to increase model skill in areas of poor performance, as well as provide a map of regions where ROMS predictions work well and where they will not provide a significant increase in navigational accuracy.

### References

- J. Binney, A. Krause, and G. S. Sukhatme, (2010). "Informative path planning for an autonomous underwater vehicle," In IEEE International Conference on Robotics and Automation.
- J. Das, K. Rajan, S. Frolov, J. Ryan, F. Py, D. A. Caron, and G. S. Sukhatme, (2010). "Towards marine bloom trajectory prediction for AUV mission planning," In IEEE International Conference on Robotics and Automation.
- A. Pereira, H. Heidarsson, D. Caron, B. Jones, and G. Sukhatme, (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), July.

R. N. Smith, Y. Chao, B. H. Jones, D. A. Caron, P. P. Li, and G. S. Sukhatme, (2009a). "Trajectory design for autonomous underwater vehicles based on ocean model predictions for feature tracking," In Proceedings of The 7th International Conference on Field and Service Robotics, (Cambridge, MA), July.

R. N. Smith, J. Das, H. Heidarsson, A. Pereira, D. A. Caron, B. H. Jones, and G. S. Sukhatme, (2009b). "Implementation of an embedded sensor network for the coordination of slocum gliders for coastal monitoring and observation," In Workshop on UnderWater Networks; ACM SenSys 2009, (Berkeley, CA), November.

R. N. Smith, J. Das, H. Heidarsson, A. Pereira, I. Cetinić, L. Darjany, M. Eve Garneau, M. D. Howard, C. Oberg, M. Ragan, A. Schnetzer, E. Seubert, E. C. Smith, B. A. Stauffer, G. Toro-Farmer, D. A. Caron, B. H. Jones, and G. S. Sukhatme, (2010a). "USC CINAPS Builds Bridges: Observing and monitoring the Southern California Bight," IEEE Robotics and Automation Magazine, Special Issue on Marine Robotics Systems, March.

R. N. Smith, Y. Chao, P.P. Li, D.A. Caron, B.H. Jones and G.S. Sukhatme, (2010b). "Planning and implementing trajectories for autonomous underwater vehicles to track evolving ocean processes based on predictions from a regional ocean model". Submitted to the International Journal of Robotics Research.

R. N. Smith, A. Pereira, Y. Chao, P.P. Li, D.A. Caron, B.H. Jones and G.S. Sukhatme, (2010c). "Autonomous underwater vehicle trajectory design coupled with predictive ocean models: A case study". In Proceedings of the IEEE International Conference on Robotics and Automation, May.

## MAS 02 A Communications Framework for the Cost-effective Operation of Slocum Gliders in Coastal Regions

### MAS 02.1 Overview

Slocum gliders have been gaining popularity as a sampling platform that provides oceanographers and marine biologists with a means to monitor the oceans much more easily over time-scales of several weeks to months without the need of having an accompanying research vessel. This results in significant savings in terms of man-power as well as a very economical way of increasing the coverage area by using multiple Slocum gliders simultaneously which provides much richer data-sets. Intelligent autonomy in these platforms makes them much easy to deploy, operate and retrieve.

This becomes even more pertinent in a coastal scenario where a coastal observatory may be able to employ a small fleet of gliders to monitor the coast almost continuously over a long period of time. In such situations, a significant economic factor turns out to be the communication cost over Iridium which makes up a large portion of the nominal operational cost of the vehicle. In a region such as the Southern California Bight, we felt we could utilize a network of base-stations at several shore-sites to cover a significant part of the operational area, which would result in a reduction in cost of the communication costs as well, time at the surface of the gliders as well as the communication bandwidth, given the vehicle's surface at appropriate locations.

To address these problems, we developed a coastal communication system ([1], [2], [3], [4]), which provides a fast, safe, affordable way of communicating with Teledyne Webb gliders through a RF-modem base-station network installed in the Southern California region. Each base-station installed at an elevation, has a computer connected to the internet and uses a custom light-weight communication protocol to communicate with surfaced Slocum gliders.

Using this method we measure a 6x improvement in file transfer speed, compared to Iridium [1]. This can be further increased upto 24x if compression is performed before transmission of typical data files. We determined experimentally that the feasible communication range using this method is approximately 12 km. While we were able to obtain links between the vehicle and the shore upto 12 km we noted that in some instances the performance of transmission could fare poorly even when the vehicle is nominally well within operating range. We hypothesize that sea state might be a factor that affects range since the glider's antenna is positioned only 30cm above the sea-surface and local waves have an increased effect in obstructing its line-of-sight to base-stations at distances further away from the base-station (See Fig.1). In order to study this, we installed a 3-axis accelerometer in one of our gliders, and collected data from it while transmitting test sequences to a base-station. The ultimate goal is to develop an intelligent system of communication that can estimate relevant wave parameters using the accelerometer and use these to determine the regime for communication on the vehicle. In this report, we discuss some results from a recent month-long field deployment in the Southern California Bight region with our system and compare wave parameters measured by our method with reference values from a moored NDBC buoy in the region.

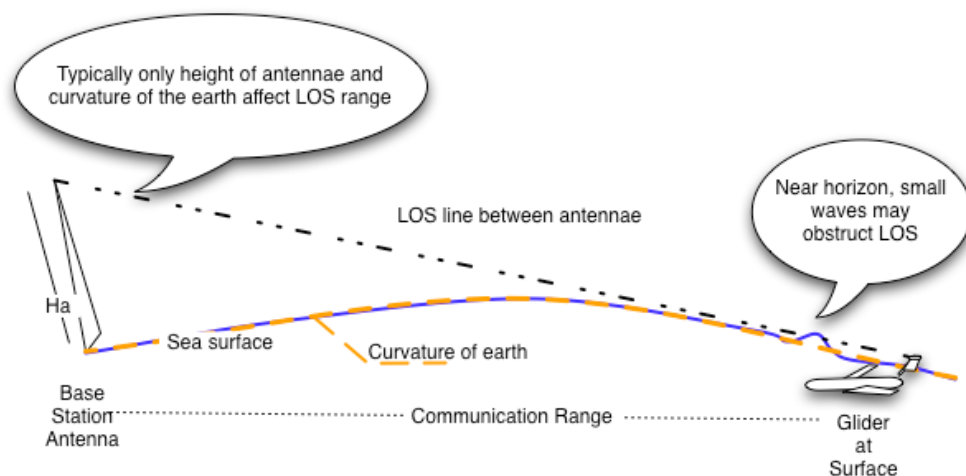


Figure 1. Depiction of glider shore Line-of-Sight (LOS) obstruction due to waves. (Not to scale).

### MAS 02.2 Approach

Unlike NDBC buoys which generally rely on 20-minute time series at 2Hz, we collect shorter accelerometer time series during glider surfacings which range between 3 to 10 minutes at a sampling rate of 50Hz. The glider experiences roll, pitch and yaw at the surface due to the hydrodynamic interactions between itself and the water. We describe an algorithm for processing accelerometer time series to obtain accelerations in the heave axis and computing wave parameters. We compare these estimates with those from the nearest buoy, which is approximately 25km away.

Time series data from the accelerometer is low-pass filtered by smoothing. After tilt-compensation using the inclinometer compass on the glider we obtain heave data (resolved acceleration along the z-axis), which we segment into lengths of 256 samples. We apply a Hanning window for spectral leakage reduction. Next, we remove the mean for each segment and then compute the FFT for each segment from which we get the acceleration spectra as described in [2]. The wave parameters computed from the acceleration spectra are significant wave height ( $H_s$ ), peak time period ( $T_p$ ) and average time period ( $T_{avg}$ ).

We find that the wave parameters obtained using the accelerometer on the glider show a similar trend to those gathered from the nearest data buoy. We trained a linear perceptron classifier on a subset of the data collected at surfacings and tried to assess how well it could be used to predict the communication quality.

The communication metric used is link percentage - defined as the ratio of the time to a persistent communication link was present with the total time at the surface. We notice from our experiments that values >90% are good for file transfers.

Training data-set	Range < 8km Link < 30%	Range < 8km Link > 30%	Range < 8km Link > 90%
A	97.25%	97.83%	96.35%
B	98.62%	97.14%	98.36%
C	98.20%	97.35%	97.72%

Table 1. Classifier performance for Ranges under 8km

Training data-set	Range > 6km Link < 30%	Range > 6km Link > 30%	Range > 6km Link > 90%
A	91.14%	90.91%	91.19%
B	90.97%	91.33%	89.40%
C	94.39%	93.62%	93.79%

Table 2. Classifier performance for Ranges greater than 6km

Training data-set	Range > 8km Link < 30%	Range > 8km Link > 30%	Range > 8km Link > 90%
A	90.62%	91.99%	90.43%
B	91.99%	91.34%	89.67%
C	94.56%	96.05%	91.67%

Table 3. Classifier performance for Ranges greater than 8km

### MAS 02.3 System(s) Description and/or Experiments

We performed two month-long deployments in the Southern California Bight region during May 2009 and September 2009. During the September 2009 trials, we collected accelerometer data while transmitting known sequences to the base-stations. The gliders were provided with missions that allowed us to collect measurements of communication performance in a region of up to 20 km from the base-station at Catalina island. We compared the extracted wave-

parameters from the accelerations on the glider to those from the nearest NDBC buoy which is located about 20km from the region where the glider surfaced.

#### MAS 02.4 Accomplishments

We find that wave parameters are useful in estimating expected quality of communication. The sea states were generally calm during the time our data was collected, with significant wave heights between 0.4m and 1.7m, and peak time periods ranging between 2.86 sec and 22.22 sec. By knowing how well we can expect to communicate over the freewave radio based on the sea-state the time at the surface can be better utilized. If the sea-state indicates that the expected communication quality will be poor, file transmissions can be postponed to a future surfacing. Accelerometer time-series from an AUV at the surface can be utilized to predict the expected quality of communication with shore stations. While range plays a major role in predicting the expected communication quality Table 2. shows that for ranges under 8km, the classifier that relies on acceleration information is more accurate in predicting when the communication quality will be very poor. This provides evidence supporting the hypothesis that sea-state does play a role in determining the quality of communication.




Figure 2: Link-state obtained from the September 2009 glider deployment. While the link state is generally very good within 8km, we find several surfacing locations where the link-state was very poor nearer the base-station on Catalina Island

### **MAS 02.5 Future Directions**

We are currently working on using our classifier to make communication decisions on when to transmit files based upon wave-parameters. We also want to vary communication parameters such as RTT (Retry-Timeout-Timers), MTU (Maximum Transmission Unit) sizes and so on. During the next few months we are doing field trials that will help us test if these ideas are practically feasible in adapting communication protocols to improve communication between gliders and shore stations. Later, we will use our communication network to drive glider re-tasking experiments which will allow multiple robot teams to be used.

### **References**

- [1] Arvind Pereira and Gaurav Sukhatme. Estimation of wave parameters from accelerometer data to aid AUV-shore communication, 2010. Submitted to Oceans 2010.
- [2] Ronald Brown, James Baker, and Jerry McCall. Nondirectional and directional wave data analysis procedures. Technical Report 01, National Data Buoy Center, NOAA, Stennis Space Center, January 1996.
- [3] USC CINAPS. <http://cinaps.usc.edu>.
- [4] Arvind A. Pereira, Hordur Heidarsson, Carl Oberg, David A. Caron, Burton H. Jones, and Gaurav S. Sukhatme. A communication frame- work for cost-effective operation of auvs in coastal regions. In The 7th International Conference on Field and Service Robots, Cambridge, Massachusetts, Jul 2009.

## MAS 03 Marine Bloom Trajectory Prediction for AUV Mission Planning

### MAS 03.1 Overview

Observing large scale oceanographic features has often been difficult. For example, Intermediate Nepheloid Layers (INL) are near-coastal fluid sheets of suspended particulates with large horizontal extent (in kilometers) and small vertical extent (in 1-10 meters). Harmful Algal Blooms (the focus of this paper) are likewise spread over kilometers with primary productivity driving its ecology within the top 5-10 meters. Traditional approaches using ship-based measurements for observing such dynamic and episodic phenomenon have proven to be ineffective given evolving biological state, the need to measure various properties across the spatial extent of such phenomenon, and most of all in dealing with logistical details centered on manned ships on fixed schedules.

More recently Autonomous Underwater Vehicles (AUVs) have shown to be more cost-effective, have demonstrated increased persistent presence, and with a suitable sensor payload have been able to systematically observe large scale phenomenon at requisite scales of variability of biogeochemical processes. Yet, such mobile robotic assets have often been ineffective in resolving the spatiotemporal characteristics to effectively sample and observe. Additionally while satellite observations have proven to be helpful, they are constrained either by cloud cover, the lack of data beyond a meter or two of the sea surface (when process driven phenomenon can and are often within the euphotic zone up to 150m in depth) and logistical issues with the time lag between observations and processed data sets available for use. Mooring data suffers from a spatial sparseness even if in-situ measurements are accurate and available in realtime. Our goal and the subject of this work is to make use of the data available from a range of sources including ocean models, remote sensing satellites, moorings, and on-shore instruments to make predictions of the trajectory of a patch of water. Our experiments target the Monterey Bay which is not only one of the most biologically diverse bodies of waters but the northeast bay frequently experiences extreme "red-tide" blooms, making it an ideal location for bloom studies. In this paper we analyze the result of advecting hotspots from blooms that occurred between October 2007 and October 2008 at the Monterey Bay, and show an example of how such predictions can be used to plan survey missions for AUVs.




Figure 1. NIR-G-B composite image of NE Monterey Bay on August 26, 2004, when an extremely dense red tide bloom was present.

### MAS 03.2 Approach

The dynamics of an algal bloom can be primarily described by three factors: advection, diffusion, and bloom ecology. Advection is the component of the transport that is due to the effect of external forcing (ocean current). Diffusion results from movement of particles along concentration gradients. Lastly, since the phenomenon is biological, its growth and decay is governed by bloom ecology. To predict the dynamics of the bloom for long timeframes (e.g., months) would require data for all three factors described above. However, for a strong coherent bloom already in progress, predictions can be made for a period of up to a week based on the external forcing and diffusion. Ocean observing systems such as SCCOOS and CeNCOOS provide near-realtime data on various aspects of the physical seastate such as temperature, current, salinity and sea surface height. Given access to these data, we ignore ecology and focus on the advective effect of ocean currents on blooms.

Our system computes prediction of bloom trajectories in two steps,

- patches of bloom hotspots are detected from remote sensing data
- detected patches are advected using surface current data from HF Radar stations.

### MAS 03.3 System Description

#### *Data Sources*

Remote Sensing Data. Remote sensing satellite data provides a synoptic view of the oceans, enabling bloom detection by use of proxy measurements, such as ocean color and emitted radiance due to fluorescence. One example is the MODIS (MODerate resolution Imaging Spectroradiometer) instrument on NASA's Terra and Aqua Earth-orbiting spacecraft. Both Terra and Aqua view the entire Earth's surface every one to two days. One of the data products from the Aqua MODIS is Fluorescence Line Height (FLH), a relative measure of the amount of radiance leaving the sea surface in the chlorophyll fluorescence emission band (685 nm), at a resolution of 1 km. FLH is a recognized proxy for chlorophyll concentration in the upper column of water. Another example is the Sea-viewing

Wide Field-of-view sensor (SeaWiFS) and the chlorophyll-a data product. Fig. 2 shows a strong bloom in the Monterey Bay that occurred between September and October 2002. MODIS data in particular, are available in hierarchical Data Format (HDF) where the FLH data product from each MODIS image can be read into a 500 x 667 matrix.

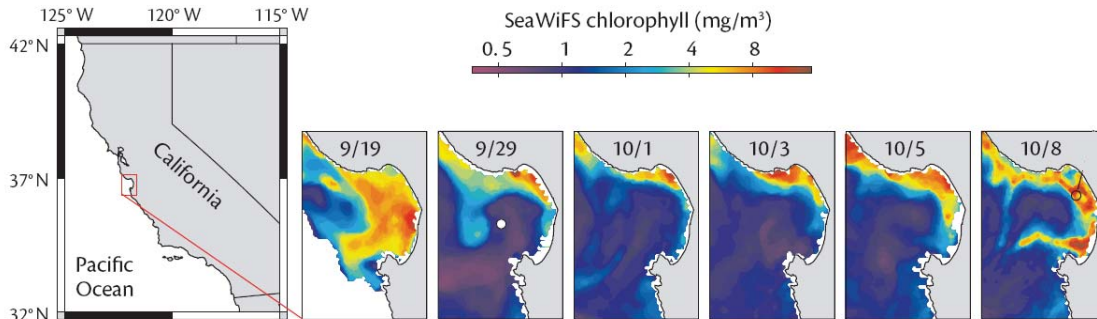


Figure 2. SeaWiFS images showing a red-tide bloom at the Monterey Bay between September and October 2002.

**Algorithm 1: Hotspot detection and advection**

```

1 Input: MODIS images  $M_1$  and  $M_2$ 
2 Time period  $T = \text{timestamp}(M_2 - M_1)$ 
3  $t = t_{M_1}$ 
4  $I_{\text{thresholded}} = \text{threshold}(M_1, F)$ 
5  $H = \text{connectedSegments}(I_{\text{thresholded}})$ 
6  $h$  hotspots,  $H = \{h_1, h_2, \dots, h_k\}$ 
7 foreach HotSpot  $h_i$  do
8   sample points for HotSpot  $h_i$ ,
9    $P_i = \text{resampleHotspot}(h_i, \text{resolution})$ 
10   $P_i = \{p_1, p_2, \dots, p_N\}$ 
11  foreach sample point  $p_j = [x, y, t]$  do
12    while  $t < t_{m2}$  do
13       $R_{p,t} = [u, v]^T$ 
14       $p_{t+\Delta t} = p_t + R_{p,t} \Delta t$ 
15       $\sigma_{p_{t+1}} = \sigma_{p_t} + \sigma_{R_t} \Delta t$ 
16       $t = t + \Delta t$ 
17    end
18  end
19 end

```

HOTSPOT PROJECTION TEST CASES		
Period(days)	Case	Evaluation Rating (max 5)
5	03/20/2008	3
	06/02/2008	4
4	10/12/2008	4
	10/09/2008	4
3	10/12/2007	4.5
	02/15/2008	3
	03/18/2008	2
	03/25/2008	2
	04/10/2008	2
	04/26/2008	2
	05/21/2008	3
	10/10/2008	3.5
	10/12/2008	3
	10/14/2008	5
1	10/24/2007	4
	10/09/2008	2

Table 1

**High Frequency Radar Data (HF Radar).** External forcing by wind and water currents is a dominating factor in bloom transport, providing reasonable estimates of the trajectory of hotspots over a period of few days using only external forcing. We used surface current data obtained from HF Radar stations maintained by CeNCOOS, which provides radial ocean surface current information in near-real time. With data from multiple stations, the velocity components of the surface current is computed. The data must be filtered, interpolated and extrapolated. In our experiments, we obtained Open-Boundary Modal Analysis (OMA) interpolated radar data for the period October 2007 - October 2008 (see Fig. 5 for an example).

(1) Hotspot detection: We define a bloom hotspot as the region in a MODIS image where the pixels have FLH values greater than a specified threshold. This is a proxy for a region of intense biological activity. Since advecting each pixel within a patch over multiple time steps is computationally expensive, we choose a sparser representation of a patch. These representative points may be from the convex hull, lie on the boundary, are from the interior, or are the centroid of the patch.

(2) Hotspot advection: Hourly HF Radar data is available at a 2 Km resolution in a gridded format. We interpolate the HF Radar surface current estimates so that for any location and time  $t$ , we have the surface current velocity. Each sample-point within a patch is projected using the interpolated HF Radar data to obtain the new location.

We performed advection of hotspots on a dataset of MODIS FLH images captured between September 2007 and November 2008 when both MODIS and HF Radar data were available. After rejecting unusable images (often due to cloud cover), we selected those that displayed hotspots of considerable intensity. We were also interested in studying the quality of advection for different time periods. From the test cases, we identified between 1 and 5 day periods between MODIS images to ground truth our projections. In total, we selected 16 test cases spanning the above period. The resulting projections, shown in Table I, were evaluated qualitatively by an oceanographer. We observed good predictions for stronger blooms, specifically those in the Fall. The advection of bloom patches fail to predict blooms that are in the initial stages of growth. However, the projections were good for blooms that were well developed and of high intensity. Fig. 3 shows examples of two-day projection from the evaluated test cases.

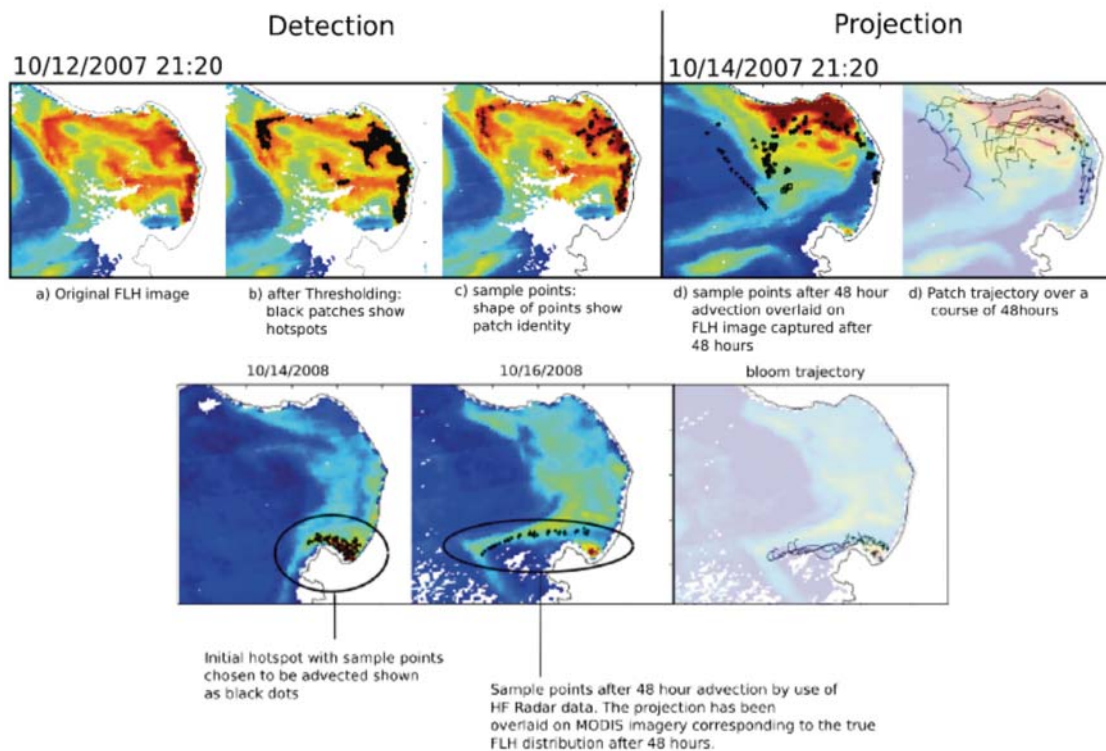


Figure 3. Projection of blooms from October 2007 and October 2008.

#### *An application in AUV sampling*

Our aim next is to choose an appropriate survey area given the prediction of a hotspot trajectory. Our objective is to project the blooms trajectory and then place a viable 'lawnmower' AUV survey to ensure coverage. A compelling scenario for such an application is illustrated in Fig. 6. We performed this step by starting with a survey area of desired dimension, and using a nested-grid approach to determine the best location and orientation for the survey that maximizes the contained bloom intensity (captured by the advected FLH values)

#### **MAS 03.4 Accomplishments**

The work resulted in a publication in IEEE International Conference on Robotics and Automation (ICRA), 2010.

#### **MAS 03.5 Future Directions**

The presented work ignored bloom ecology and diffusion for the predictions. The current focus is on introducing predictions from statistical bloom models into our system. Specifically, we are working towards model-based autonomy that integrates output from physical ocean models and biological bloom prediction models to deliberative planners to allow AUVs to perform missions with multiple science goals..

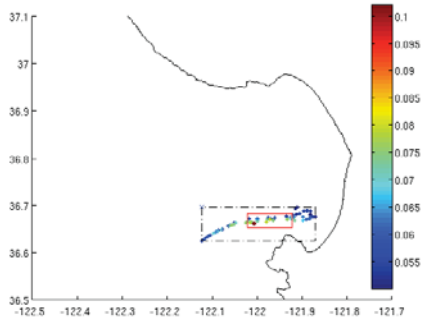


Figure 4. Plot showing the survey rectangle for the October 2008 bloom

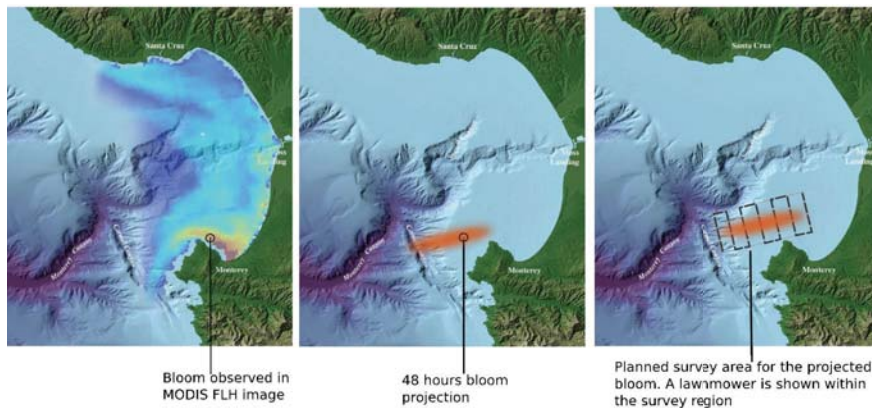


Figure 5. A hypothetical bloom onset (A), projection (B), and a viable AUV survey to capture the bloom's spatial extent.

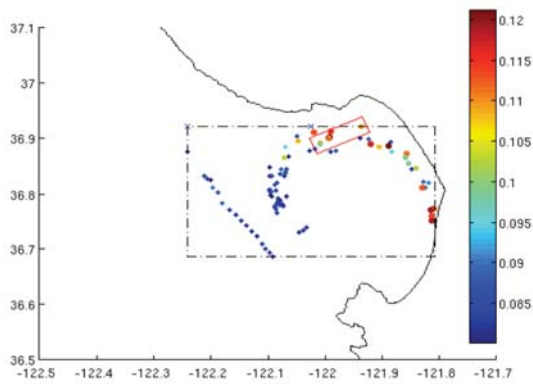


Figure 6. Rectangular area of known size that maximizes the contained FLH intensity of the bloom case 1 (10/12/2007). The nested grid approach was used to search for the location and orientation of the rectangle that maximized the gain (FLH intensity).

## MAS 04 Calibration and Observability

### MAS 04.1 Overview

Cameras are versatile sensors that are able to provide dense information about the local environment. When combined with inertial sensing, vision offers a very robust solution for mobile sensor pose estimation. However, to correctly fuse the data from a camera and an inertial measurement unit (IMU) into a single reference frame, the 6-DOF transform between the sensors must be accurately known. Calibration of this transform is time consuming and usually requires additional equipment. We have developed an algorithm, based on the unscented Kalman filter (UKF), for self-calibration of the transform between a camera and an IMU. Self-calibration refers to the process of using imperfect (noisy) measurements from the sensor themselves to improve our estimates of related system parameters.

Visual-inertial sensor self-calibration is challenging for several unique reasons. IMU measurements, i.e. the outputs of three orthogonal angular rate gyroscopes and three orthogonal accelerometers, can be integrated to determine the change in sensor pose over time. However, all inertial sensors, and particularly the low-cost units suitable for wide deployment, are subject to drift. IMU accelerometers also sense the force of gravity in addition to forces which accelerate the platform, gravity will normally dominate other measured accelerations. If the orientation of the IMU with respect to gravity is unknown or is misestimated, then the integrated sensor pose will diverge rapidly from the true pose. While measurements derived from camera images are immune to drift (as long as the same landmarks remain within the camera's field of view), cameras are bearing-only sensors which are not able to recover the true scale of a scene.

The goal of our work is to demonstrate that it is possible to accurately calibrate the relative pose of these sensors in the field and without any additional calibration apparatus. We foresee broad applications for this research in areas related to environment mapping and sampling by mobile sensor nodes.

### MAS 04.2 Approach

We recently extend our work on camera-IMU calibration to omnidirectional sensors and to full self-calibration, i.e. calibration without the use of a known calibration object. We formulate calibration as a filtering problem, and use the UKF to estimate the 6-DOF camera-IMU transform, the time-varying IMU biases, the local gravity vector and the metric scene structure. That is, we simultaneously calibrate, localize the camera-IMU platform and build a sparse map. Further, our approach does not require any prior knowledge about the environment in which a mobile node or robot is operating.

Our filtering algorithm is based on a differential geometric analysis of the observability of camera-IMU self-calibration. This analysis shows that we are able to self-calibrate as long as the IMU measures two non-zero angular rotation rates and two non-zero linear accelerations (i.e. along at least two axes).

### MAS 04.3 System Description and Experiments

Our current test system includes a black and white Flea FireWire camera from Point Grey Research (640 x 480 pixel resolution), coupled to a NetVision 360 omnidirectional lens manufactured by Remote Reality Inc. (shown in Figure 1). Images are typically captured at a rate of 3.75 to 15 Hz. Our IMU is a MicroStrain 3DM-GX3 unit, which provides three-axis angular rate and linear acceleration measurements at 100 Hz. Scale and axis non-orthogonality effects are compensated for internally by the IMU.

For our latest omnidirectional calibration experiments, we attached the sensors to a Directed Perception PTU-D46 pan-tilt unit, which was mounted on a Pioneer 2-AT mobile robot. The use of a pan-tilt unit made it possible to excite two degrees of rotational and two degrees of translational freedom while the robot moved on a smooth surface planar surface.



Figure 1. Pioneer 2-AT robot equipped for calibration experiments. The Pioneer carries a NetVision 360 omnidirectional camera and MicroStrain 3DM-GX3 IMU.

**MAS 04.4 Accomplishments**

Our experiments have shown that it is possible to accurately self-calibrate the 6-DOF transform between an omnidirectional camera and an IMU online and without any additional hardware. This is possible even when low-cost, solid state inertial sensors are employed.

In the past year, we reported these results in a conference paper (which appeared in the 2009 IEEE International Symposium on Computational Intelligence in Robotics and Automation), and in an article submitted to the International Journal of Robotics Research (currently in review).

**MAS 04.5 Future Directions**

As future work, we are currently exploring the possibility of self-calibrating the camera lens intrinsic and distortion parameters along with the relative pose of the sensors. This would enable complete camera-IMU self-calibration, and allow the sensor system to operate for significant durations without manual recalibration. We are also planning a series of long-distance motion estimation studies, in order to better quantify the effect of calibration on performance.

## MAS 05 Informative Path Planning

### MAS 05.1 Overview

My goal is to create path planning algorithms that give mobile robots a high degree of autonomy. Many projects within CENS involve environmental monitoring. Static sensors can provide high temporal resolution measurements, but lack spatial resolution unless a very large number are used. Because of this, it is advantageous to use robots that can traverse the environment and provide samples on demand. In order to best take advantage of the mobile sensing capability that robots provide, it is necessary to come up with a path planning algorithm that can choose paths which gather the most useful information.

Over the past year, I have focused primarily on path planning for underwater gliders, which our lab has two of. The gliders have sensors which can measure the temperature and salinity of the water, among other things. They move slowly through the water, but are capable of remaining at sea for weeks at a time. My work has focused on creating and testing informative path planning.



Figure 1: One of our underwater gliders resting at the surface.

### MAS 05.2 Approach

We take an Informative Path Planning approach to the robot path planning problem. In this approach, we aim to choose measurements best describe a scalar field of interest (e.g. temperature or salinity). We assume that we have some model for the underlying scalar field which allows us to compute the covariance and entropy of the field given a hypothetical set of sample locations. Using this model of the underlying field, our goal is to maximize the mutual information of the unsensed locations with respect to the sensed locations.

We specify the path planning problem as one of finding a path on a discrete graph. Each point in the graph represents a possible waypoint for the glider, and each edge represents a possible path leg between two waypoints. The weight of each edge represents the cost (in time or power) of traversing it. Our goal is to find a path from a node  $s$  to a node  $t$ , the total cost of which is less than some budget  $B$ , which maximizes the mutual information between sensed and unsensed locations.

Because the number of possible paths is exponential in the maximum path length, exhaustive search is not practical even for relatively small graphs. Fortunately, mutual information is submodular, which allows a more efficient approximate algorithm. Submodularity is basically the property of diminishing returns: once many measurements have already been taken, taking more measurements provides less new information.

Because of the submodularity of mutual information, we are able to use a modified version of Chekuri and Pal's recursive-greedy approximation algorithm. This algorithm provides a logarithmic approximation to the submodular orienteering problem in sub-exponential time.

### MAS 05.3 Experiments

We tested our path planning algorithms using both simulated data and data collected in the field.

#### Experiments Using Real Data

To test our algorithm, we used salinity data collected by underwater gliders. We used a Gaussian process to model the data, and restricted the space of possible paths to two dimensions. Fig. 2 shows one resulting path, and the pilot data used.




Figure 2: Pilot data (x's) and path planned by our algorithm (dotted line).

#### *Experiments Using Simulated Data*

Because we do not have ground truth values for each location, there is a limit to the kinds of testing that can be done on real data. To do more thorough testing, we use data from the Regional Ocean Modeling System (ROMS) which is maintained by JPL. This system incorporates large amounts of sensor data to create estimates of temperature, salinity, and currents for the ocean off of the Southern California coast.

#### **MAS 05.4 Accomplishments**

During the reporting period, we have done several field tests to collect data for testing our algorithm. We have also done tests in simulation, and we have extended the algorithm itself to handle new and useful cases. The initial results of this work is given in our ICRA 2010 paper.

#### **MAS 05.5 Future Directions**

Over the next year, we would like to expand the usefulness of our algorithm by incorporating ocean current estimates to get more accurate travel times for the glider. We also plan on extending our algorithm to work on other kinds of mobile robots.

## MAS 06 Imagers as Sensors

### MAS 06.1 Overview

There exist many biological sensing applications where direct measurement is either impossible, extremely invasive, or extremely time consuming. For example, measuring the presence/absence of birds at a feeder station currently requires a human to watch a camera pointed at the feeder, identifying when birds arrive and leave. Similarly, measuring CO<sub>2</sub> flux from a plant requires placing the plant inside a growth chamber, destructively modifying the environment. We propose using imagers as biological sensors by constructing a procedure that uses images to obtain approximate measurements of these phenomena. This procedure, composed of state-of-the-art computer vision, image processing, and statistical learning algorithms, will be evaluated in the context of a specific application and shown to be general through multiple instantiations. Through application, it has been found that many of these algorithms make unacceptable assumptions about their input. Providing accurate data to biologists and ecologists, though the appropriate modification of these algorithms, is the ultimate goal of this work.

In addition to our previous goal, we are now also interested in providing a system available to the general public that can segment streaming video of interesting events.

### MAS 06.2 Approach

**Animal Monitoring:** Manually monitoring the environment for the presence of animals is limited by the amount of human time that can be devoted to it. Automatically detecting and segmenting out animals from a natural scene would ease the burden of the biologist and allow observations to scale beyond what a biologist has time to view. Unfortunately, environmental monitoring applications present a challenge to current background subtraction algorithms that analyze the temporal variability of pixel intensities, because of the complex texture and motion of the scene. They also present a challenge to segmentation algorithms that compare intensity or color distributions between the foreground and the background in each image independently, because objects of interest such as animals have adapted to blend in.

**Pollinator Counting:** We would like to detect the presence of somewhat rare, novel objects (e.g., bees) near specific, user-defined regions of interest (e.g., flowers). We leverage the highly correlated nature of temporally adjacent frames within the video sequence to both reinforce potential detections and discount potential mis-detections of novel objects. Sequences of frames thought to contain such novel objects (the target) would then be provided to the domain scientist for further study. Such a system allows domain scientists to manually analyze orders-of-magnitude less video data in this summarized form as compared with unaided, human analysis of the complete video.

### MAS 06.3 System(s) Description and/or Experiments

We developed multi-stage procedures for each research area to process the incoming video imagery and produce occupancy data that can be consumed by scientists. The procedures are designed to be tailored to the specific application of interest, though the procedures themselves are general. Further, we attempt to minimize the user input required to extract the biological signal of interest.

**Animal Monitoring:** Our approach is to first localize birds through exploiting the fixed view deployment scenario, and cluster detections based on the similarity of appearance and spatial temporal coherence. We introduce a discrepancy measure that has several useful properties for multi-view categorization in the context of surveillance and monitoring applications. First, it is not biased towards dominant views in either the training or the test sets. This is particularly important because we can not guarantee a uniform distribution of collected views as a multi-camera approach would be able to.

**Pollinator Counting:** The procedure roughly consists of three independent pieces. For each frame, we first localize the ROI as specified by the user. This effectively discards a large fraction of each frame that is uninteresting, by definition, and would only serve to confuse further processing. Next, we detect and localize the target as it occludes the ROI. Finally, using the match value and location of a potential match produced by the target detection, we track the target's motion over time, leveraging the inherent temporal correlation between frames to discard false detections.

### MAS 06.4 Accomplishments

#### *Animal Monitoring*

We found that our new background subtraction improved the performance relative to other background subtraction methods. A sample foreground/background segmentation is shown in Figure 1. Our proposed multi-view categorization method also outperformed equivalent single-view approaches. We show some sample clusters derived from an unsupervised method using our discrepancy metric in Figure 2.




Figure 1. Extracting regions from a single frame. From left to right: (a) the original frame, (b) a pixel wise comparison against the background model where white is very different from the background, and (c) the highlighted regions used for object and category classification. The pink labels beneath the regions indicate the detection threshold needed for the region to be considered foreground.

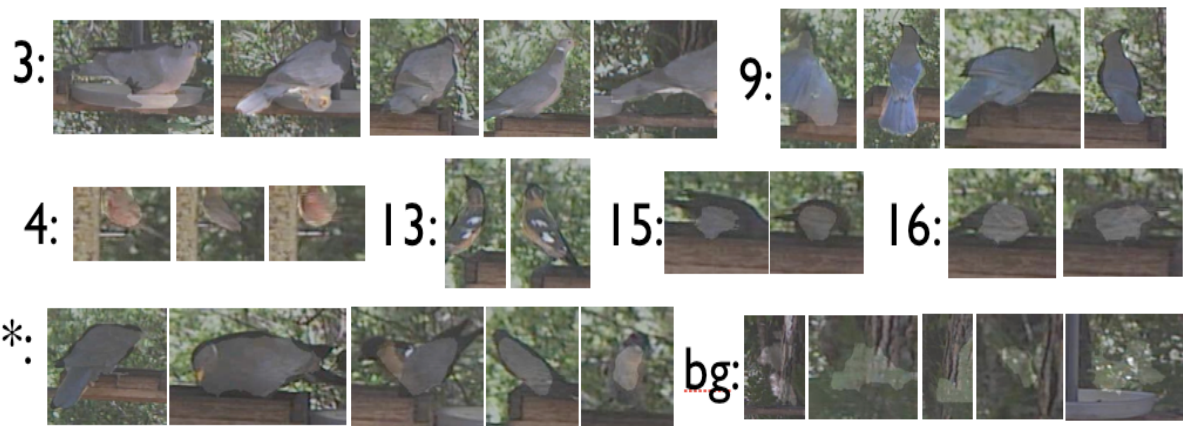


Figure 2. The clusters labeled “3, 4, 9, 13, 15, 16” are a few of the categories that were cleanly segmented from the image sequence. The “\*” category is one where our approach grouped multiple species into a single cluster. The middle image, where one bird is partially occluded by another bird, is one of the reason for this confusion. A region contains the color features from two different bird categories resulted in both species being combined into one category. The “bg” category is an example of one of the clustered background objects.

### Pollinator Counting

In Figure 3, we illustrate the challenge of detecting pollinators from images. Not only is the pollinator a small part of a complex and cluttered scene, the flower on which the pollinator rests sways in the wind throughout the image sequence. Using template alignment, initiated through user input, we can stabilize the flower and then use a background subtraction technique to detect the pollinator shown in Figure 4.




Figure 3. Against a complex and cluttered background (a), even a human observer would have trouble identifying the target. However, when restricting our view to only a important region of interest (ROI) (b), the target stands out more visibly.



(a) ROI with target

(b) Difference

Figure 4. Background subtraction has more difficulty modeling cluttered natural scenery (a) where light flecks and other transient effects cause increased noise in the difference image (b).

### **MAS 06.5 Future Directions**

We will continue the animal monitoring work by investigating methods to improve categorization and clustering of animals. Also, we will develop a working system to investigate the stability of the algorithms we have proposed, and validate their usefulness.

## MAS 07 Multi-Class SVM for Forestry Classification

### MAS 07.1 Overview

In this project, we propose a method for classifying the vegetation types in an aerial Color Infra-Red (CIR) image. Different vegetation types do not only differ in color, but also in texture. We study the use of four Haralick features (energy, contrast, entropy, and homogeneity) for texture analysis, and then perform the classification using the One-Against-All (OAA) multiclass Support Vector Machine (SVM), which is a popular supervised learning technique for classification. The choice of features (along with their corresponding parameters), the choice of the training set, and the choice of the SVM kernel highly affect the performance of the classification. The study was done on several CIR aerial images provided by the French National Forest Inventory (IFN). We will show one example on a national forest near Sedan (in France), and compare our result with the IFN map.

### MAS 07.2 Approach

We present a method to classify tree types in an aerial Color Infra-Red (CIR) image. The method consists of extracting relevant image features that enhance the distinguishability of the tree types, and then performing the classification in the feature space. The Haralick features were chosen for the feature extraction, and the Support Vector Machine (SVM) for the classification. The French National Forest Inventory (IFN) provided the image that we used. It contains two types of trees, which makes the problem a three-class classification one. This research was mainly motivated by the need of IFN to improve their forestry maps. Fig.3, taken from the IFN registry available on [1], is the IFN map for the test image (Fig.1), where it is clear that the map does not show enough details.

#### *Support Vector Machines*

The support vector machine (SVM) is one of the most popular methods for classification. It has been used in many domains and applied to different classification problems [2], [3], [4]. SVM falls in the category of supervised learning techniques, and consists of finding the optimal separating hyperplanes between the classification classes, based on training data. The training data are representational data from each class. The task is, then, to find the optimal hyperplanes separating the training data, and classifying a new datum depending on where it falls with respect to the hyperplanes. Mathematically, we seek to find a linear function  $f(x)$  that classifies the points. The popularity of SVM is due to its small generalization error, and its low computational complexity as it is formulated as a convex program. Suppose we have two classes and training data  $x_i$  in  $R^d$  with corresponding labels  $y_i$  in  $\{-1; +1\}$ , where  $i = 1, \dots, N$ . The convex optimization problem that finds the hyperplanes is the following:

$$\begin{aligned} \max_{\alpha} \quad & \sum_{i=0}^N \alpha_i - 1/2 \sum_{i=0}^N \sum_{j=0}^N \alpha_i \alpha_j y_i y_j K(x_i, x_j) \\ \text{subject to} \quad & \sum_{j=0}^N \alpha_j y_j = 0, \quad 0 \leq \alpha_i \leq C \end{aligned}$$

where  $K(\dots)$  is a kernel function satisfying Mercer's theorem, and  $C$  is a regularization constant to allow some missclassification when the classes are non-separable.

The classification function  $f(x)$ , is given by

$$f(x) = \sum_{i=0}^N \alpha_i y_i K(x_i, x) + b$$

Where  $b$  is found from the KKT conditions.

The most popular approach to build a multi-class SVM, is to combine multiple binary SVMs. In this project we use the One-Against-All (OAA) scheme, which consists of building one SVM for each class [5]. Given  $M$  classes of data, we construct  $M$  binary SVMs where the  $i$ th SVM separates class  $i$  (labeled positive) from the rest  $M-1$  classes (labeled negative), where  $i = 1, \dots, M$ . The resulting  $M$  classification functions,  $f_i(x)$ , are then combined by a max rule to form one global classification function. Given the labeling assumed, a data point  $x$  is assigned to the class corresponding to the largest classification function:

$$\text{class}(x) = \arg \max_i f_i(x)$$

### Haralick Features for Texture Extraction

Different types of trees do not only differ in color but also in texture. For texture extraction, we investigated the use of the popular Haralick features [6], which are based on the gray level co-occurrence matrix (GLCM). The GLCM represents the distribution of repeated values at a given offset  $(\Delta x, \Delta y)$ , in an image (or a window of an image) [12].

Mathematically, for an  $A \times B$  window,  $I$ , of an image with  $G$  gray levels, the GLCM  $C$  is a  $G \times G$  matrix defined by:

$$C(i, j) = \sum_{i=0}^A \sum_{j=0}^B \begin{cases} 1, & I(i, j) = i, I(i + \Delta x, j + \Delta y) = j \\ 0, & \text{otherwise} \end{cases}$$

To get a normalized distribution  $p$ , we normalize  $C$

$$p(i, j) = \frac{C}{\sum_{i=0}^G \sum_{j=0}^G C(i, j)}$$

Haralick defined 13 features in [12]. In this work, we chose to study four of them (energy, contrast, entropy, homogeneity):

$$\text{Energy} = \sum_{i=0}^G \sum_{j=0}^G p^2(i, j)$$

$$\text{Contrast} = \sum_{i=0}^G \sum_{j=0}^G (i - j)^2 p(i, j)$$

$$\text{Entropy} = - \sum_{i=0}^G \sum_{j=0}^G p(i, j) \log p(i, j)$$

$$\text{Homogeneity} = \sum_{i=0}^G \sum_{j=0}^G \frac{p(i, j)}{1 + |i - j|}$$

In addition to the Haralick features, we studied the use of the mean, the standard deviation and the normalized difference vegetation index (NDVI). The NDVI is computed from the spectral values of each pixel, and reflects the vegetation density for the pixel [2]. Finally, we chose  $G$  to be 32.

### MAS 07.3 Experiments and Results

The image has three spectral bands (NIR, red and green), and contains two types of trees (coniferous, and leafy). Our goal is to classify each pixel, to check if it belongs to the first type of trees, the second type, or neither. Thus we have a three-class classification problem, where we call the classes coniferous, leafy, and other. The other class contains the non-vegetation, the very young regeneration, and the shadows.

As mentioned above, we have a total of 7 features, which we compute for each pixel of the image. All of them, except the NDVI, are computed over a window centered at the pixel. We tested different window sizes: 5x5, 7x7, 11x11, 15x15, 17x17, 21x21 and 25x25, and we chose the 17x17 as it showed better distinguishability of textures than the other sizes. A lot of work could be done for more rigorous window size selection, but this was out of the scope of the current work. We have plans to study it more carefully in the future.

After the study of the 7 features, we decided to only use the NDVI, the mean, the standard deviation, and the entropy. The other features did not seem to discriminate well between the different classes. In addition to that, we used the spectral bands (NIR, red, and green). In summary, vectors of the form (NDVI, NIR, red, green, mean, standard deviation, entropy) were the inputs to the multi-class SVM.

For the SVM, we used the Gaussian kernel

$$K(x_i, x_j) = \exp\left(-\frac{\|x_i - x_j\|^2}{\sigma^2}\right)$$

The choice of the SVM parameters ( $\sigma$ , and  $C$ ) was done by running a series of simulations over a range of values, on the training set. We chose the ones that resulted in the best classification of the data set. We did not have a ground truth for evaluating the classification, and the evaluation was done visually. The parameters will work well for the same type of images with the same resolution. We tested the same parameters on different images of the same type

of trees, but with different resolutions and the results were not so good. The parameters needed to be changed to get a better classification on the other image. The classification result is shown in Fig. 3, where blue represents coniferous, green represents leafy, and white represents other. This result was obtained for  $\sigma = 0.1$ , and  $C = 1$ . This result shows an improvement over the IFN map shown in Fig. 2. The closed forests were very well classified; the three classes were clearly distinguished even in the overlapped regions. In the open forests, the classification problems were due on the one hand to the shadow, and on the other to the fact that the young regeneration has a similar reflectance and texture as the trees. However, in the open forest our result shows a noticeable difference between the regions where the regeneration was cleared (look at the thin white lines in Fig.1), and the ones where the clearing was not done. This differentiation is very important for forest management. We are working on removing the artifacts and making the classification smoother.




Fig. 1. The test CIR image. ©IFN.



Fig. 2. The current IFN map. ©IFN.




Fig. 3. The classification result. Blue for coniferous, green for leafy, and white for other.

#### MAS 07.4 Accomplishments

We showed in this work, that texture analysis and SVMs work well on classifying forestry type in CIR images. The initial result presented here, show that we can get good details of the forestry map. A paper about this work was published in the International Conference on Image Processing (ICIP09):

- "Multi-Class SVM for Forestry Classification." N. Hajj Chehade, JG. Boureau, C. Vidal, and J. Zerubia. In Proc. IEEE International Conference on Image Processing (ICIP), Cairo, Egypt, November 2009.

### **MAS 07.5 Future Directions**

For the coming year, we plan on extending this work in three directions:

- **Parameter learning:** As we said above, in this work the SVM parameters were chosen manually. We plan on developing a method for automatic parameter learning from the training data. We currently have two ideas for this task, the first one is to use a Bayesian framework. The second is to formulate an (convex) optimization problem for the parameters.
- **Multi-level classification:** In this problem, we are also interested in knowing the density of the forest, and whether the regeneration was cleared or not. We are planning to design a multi-level SVM for the classification. A first level would classify the type of the tree, a second level would classify the density (dense or sparse), and finally a third level would classify the regeneration.
- **Human activity classification:** We plan on trying SVMs for the classification patterns in time series medical data. The goal of the project is to classify the human activity and we believe that SVMs will be good for this task.

### **MAS 07.6 References**

[1] "Inventaire Forestier National," [www.ifn.fr](http://www.ifn.fr).

[2] O. Zammit, X. Descombes, and J. Zerubia, "Support vector machines for burnt area discrimination," Research Report 6343, INRIA, November 2007.

[3] B. Waske, G. Menz, and J.A. Benediktsson, "Fusion of support vector machines for classifying SAR and multispectral imagery from agricultural areas," Proceedings of the International Geoscience and Remote Sensing Symposium, pp. 4842–4845, July 2007.

[4] C. J. C. Burges, "A tutorial on support vector machines for pattern recognition," Data Mining and Knowledge Discovery, vol. 2, pp. 121–167, 1998.

[5] C. W. Hsu and C. J. Lin, "A comparison of methods for multiclass support vector machines," IEEE Transactions on Neural Networks, vol. 13, no. 2, pp. 415–425, 2002.

[6] R. M. Haralick, "Statistical and structural approaches to texture," Proceedings of the IEEE, vol. 67, no. 5, pp.786–804, 1979.

## MAS 08 Optimum Joint Event and Parameter Estimation in SN Based on Random Set Theory

### MAS 08.1 Overview

In SN, for dynamic events, classical joint estimation/detection/tracking multi-sensor and multi-target algorithms are often hybrids of both analytical and ad-hoc approaches at various levels. The intricacies of the resulting solutions when the number of targets and sensors may vary randomly often obscure design intuition and leave many design choices to a largely trial and error based approach. Random Set Theory is a formal generalization of the classical random variable into the random set domain.

### MAS 08.2 Approach

By treating multi-target and multi-sensor jointly, RST is able to provide a systematic framework for rigorous mathematical analysis. Because of its set theory domain, RST is able to model the randomness of transmission noise, missed detection, sensor failure, target appearance and disappearance, clutter, jammer, ambiguous measurements, and other practical artifacts within its probabilistic framework. Furthermore, a rigorous statistical framework has been developed for RST that includes concepts as: ML, Bayesian filtering, data fusion, and Creamer-Rao Bound.

### MAS 08.3 System(s) Description and/or Experiments

We initiated a proof-of-concept use of RST to jointly detect/locate/track targets successfully in a simple power aware spatially dense uniform grid WSN setting. Consider a simple model using a large number of low cost sensors placed in a uniform grid field, in which when each sensor is activated, it sends the detected signal strength of the source within its limited sensing range, using its radio to a fusion center for processing. Due to a sensor's limited sensing range, we can assume a sensor is only activated by a single source, while there may be other sources in the entire field. Most sensors are in the "sleep mode", capable of detection but not transmission. In this model, there are at least three kinds of randomness in the system. First, there is the conventional random noise in the radio transmission channel of each sensor. Second, the number of sources (i.e., targets) in the field may also be random (since non-static sources may enter and leave the field in an unknown manner). Third, the number of sensors actually able to sense and perform a successful radio report may also be random, due to battery or circuit failures, or severe propagation losses. The systematic treatment of the last two kinds of randomness, which is intrinsic to a WSN, should not be performed in an ad hoc manner, but should best be handled using RST.

### MAS 08.4 Accomplishments

Consider a WSN scenario, where we deploy 100 sensors distributed on a 100 by 100 meter square field in a uniform grid. When a target is present, it emits signal power  $P_0$  that decays  $r^{-\alpha}$ , where  $1 < \alpha \leq 2$ . Each sensor will receive this signal that gets corrupted by an AWGN with variance  $\sigma^2$ . If the received signal is above the threshold  $\tau$ , the sensor reports its reading to the fusion center. At each time instant, the fusion center jointly detects whether there is a target, and when there is one, estimates or even tracks its position. In Figure 1-a, we show the tracking of a target moving across the field when all the sensors are working properly, while in Figure 1-b, some sensors have failed to report and the tracker is not aware of their failure. The resulting tracks (shown by the x symbols) are poor. In Figure 2-a, when the fusion center is running the RST tracking system, the failed sensor information are used automatically by the RST system, and the target tracking suffers momentarily. Figure 2-b shows under the RST operated system, future tracks recover properly. In Figure 3-a, the RFS fusion center tracks a target for 11 time instants. Then it

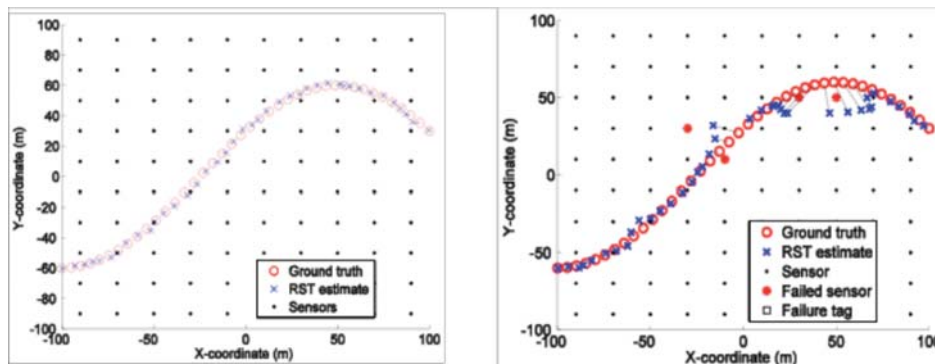


Figure 1. Left. All sensors operate properly. Right. Some sensors failed to report resulting in poor tracks.

disappeared from time 12 to 15, but a target reappears from time 16 to 31. The RST fusion center can re-initiate the tracks properly. In Figure 3-b, the RST tracker estimates only one source from time 1 to 15, zero source from time 12 to 15, and one source from time 16 to time 31.




Figure 2. Left. Two sensors failures detected by the RST fusion center. Right. Future tracks recover properly.

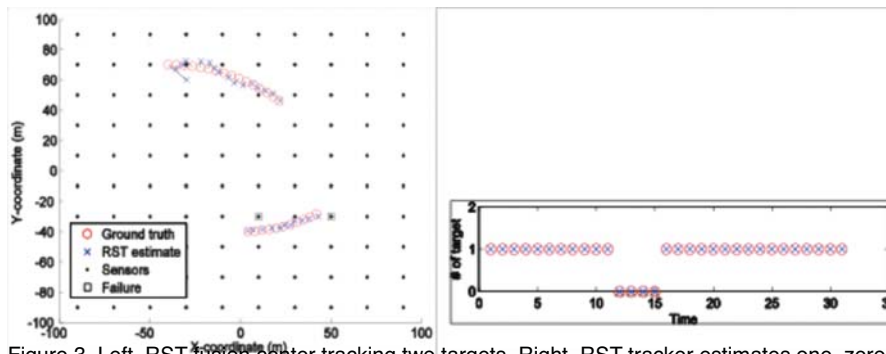


Figure 3. Left. RST fusion center tracking two targets. Right. RST tracker estimates one, zero, and then one target.

### MAS 08.5 Future Directions

Passive acoustic tracking benefits animal bio-behavioral study in replacing or enhancing human involvement in performing field data collection. Multiple simultaneous vocalizations are a common occurrence in a forest or a jungle, where many species are encountered. Given a set of nodes that are capable of producing multiple direction-of-arrivals estimates, such data needs to be combined into meaningful tracks. RST provides the mathematical probabilistic model, which is suitable for analysis and optimal fusion center synthesis. Our proposed algorithm will be tested, modified (incorporating conceptual reformulations), and retested based on real life experimentally collected data.