Exploration of the Gulf of Mexico Oil Spill with the *Sentry* Autonomous Underwater Vehicle

Michael V. Jakuba*, James C. Kinsey†, Dana R. Yoerger†, Richard Camilli†, Christopher A. Murphy†, Daniel Steinberg*, and Asher Bender*

*Australian Centre for Field Robotics
University of Sydney
Sydney, NSW 2041, Australia
{m.jakuba|d.steinberg|a.bender}@acfr.usyd.edu.au

†Deep Submergence Laboratory
Woods Hole Oceanographic Institution
Woods Hole, MA 02543, USA
{jkinsey|dyoerger|rcamilli|cmurphy}@whoi.edu

***Abstract*— We report on the use robotic assets in the investigation of subsea hydrocarbon plumes caused by the blowout, on 20 April 2010, and subsequent sinking of the *Deepwater Horizon* drilling platform in the Gulf of Mexico. We employed conventional oceanographic sampling techniques along with the *Sentry* autonomous underwater vehicle (AUV) to confirm the existence of a coherent subsea hydrocarbon plume, then to map the plume's spatial extent out to 35 km down-current from the well head, and finally to collect targeted water samples from within the plume itself for later laboratory analysis. In this paper we focus on the techniques used to coordinate sampling activities between the AUV and conventional instrumentation: geo-referenced navigation of all data, integrative visualization of multi-modal and multi-platform data, real-time telemetry, visualization, and analysis of data, and the real-time adaptation of vehicle trajectory in response. Our results demonstrate that when initial characterization is poor, limited human interaction and feedback can accelerate the study, and improve the analysis, of evolving environmental phenomena. We discuss several lessons learned, particularly as they apply to the future development of limited-interaction autonomy in subsea robotics. Using real data collected during the *Deepwater Horizon* expedition, we present simulations of semi-automated data interpretation and sampling plan adaptation comparable to the real-time actions taken by us during the expedition itself.**

Fig. 1. The Woods Hole Oceanographic Institution AUV *Sentry* on the deck of the R/V *Endeavor* between deployments. Drilling platforms and other equipment working at the *Deepwater Horizon* blowout site are visible in the background. The closest to the site that the vehicle was deployed was 3 nautical miles (5 km). Photo by D. Yoerger.

I. Introduction

On 20 April 2010, the *Deepwater Horizon* drilling rig suffered a blowout that resulted in the eventual sinking of the rig and the death of 11 personnel on board. Prior to the successful capping on 15 July 2010, oil from the damaged well head was leaking at a rate whose quantification remains contentious but undoubtedly represents one of the largest accidental releases of oil on record. The environmental impact of the oil spill depends on a number of incompletely understood characteristics of the spill including composition of the oil, its chemical evolution in the environment, the rate and total volume of oil released, and the dynamics of its spread.

Since shortly after the explosion various sources have reported the presence of subsurface plumes of oil [1], [2]. Contrary to the elementary notion that oil and water do not mix, the mixture emanating from the well head is a complex multi-phase mixture of oil and gases that interacts with the surrounding water column as it rises. Both controlled experiments [3] and historical evidence [4] suggest that some constituents of the effluent and/or minute droplets of oil will enter the water column forming a subsurface plume with little or no residual buoyancy. The composition of any subsurface plumes and the fraction of the total oil released that they represent could play a significant role in the ultimate environmental and economic impact of the spill.

In June 2010 the authors were part of a research cruise to the Gulf of Mexico funded by the United States National Science Foundation to identify and characterize any subsurface plumes associated with the *Deepwater Horizon* spill. We employed two principal sampling platforms, a conventional cable-lowered oceanographic conductivity, temperature, and depth (CTD) rosette augmented with a TETHYS in situ mass spectrometer [5] as well as several sensors specifically selected for the cruise; and the Woods Hole Oceanographic Institution's *Sentry* Autonomous Underwater Vehicle (AUV) also equipped with a TETHYS instrument as well as various other water column sensors (Fig. 1). The lowered CTD included the ability to collect water samples — crucial to determining the exact composition of the plume (Fig. 2).

This paper is organized as follows. Sec. II discusses the techniques used to coordinate sampling activities between the CTD and AUV for plume localization, plume characterization, and targeted water sampling: geo-referenced navigation of all data; integrative visualization of multi-modal and multi-platform data, real-time telemetry, visualization, and analysis of data, and the real-time adaptation of vehicle trajectory in response. Sec. III proposes a method for conducting subsea robotic survey that capitalizes on the increasing availability and bandwidth of acoustic communications for real-time human interaction with subsea assets combined with modern machine learning techniques for dimensionality reduction and data pre-processing. The method aims to enable human operators to focus on high-level data interpretation and mission objective formulation when adapting sampling plans. We demonstrate the method on data collected during the *Deepwater Horizon* expedition using a data-denial methodology to simulate real-time vehicle trajectory adaptation and provide a qualitative assessment of the results relative to actions taken by us to adapt vehicle trajectory during the actual expedition. We conclude with a discussion of lessons learned.

II. Subsurface Plume Localization at the *Deepwater Horizon* Blowout Site

Our ultimate objective was to collect targeted water samples from within any subsurface plumes for later analysis in shoreside laboratories. To accomplish this we first had to confirm the existence of plumes by locating and mapping them. Our approach capitalized on the strengths of our two sampling platforms, while respecting the constraints imposed by the sensors on board. The lowered CTD was used to initially locate the plume and then to characterise its vertical structure. The AUV provided a complementary horizontal perspective.

Fig. 3 shows the northeast corner of the Gulf of Mexico and the location of the *Deepwater Horizon* site off the Louisiana coast. With the exception of background water profile measurements all subsurface operations during the 10 days we spent on station took place within the area indicated. We conducted 23 CTD lowerings including 3 extended deployments in which the instrument package was towed slowly while undulating within a prescribed depth interval (a procedure known in the oceanographic community as a tow-yo). The CTD data identified a deep subsurface plume tending to the WSW of the *Deepwater Horizon* site and centered at a depth of 1100 m. *Sentry* dived 3 times, covering approximately 240 kilometers at depths between 1000 m and 1300 m. Two of these dives,

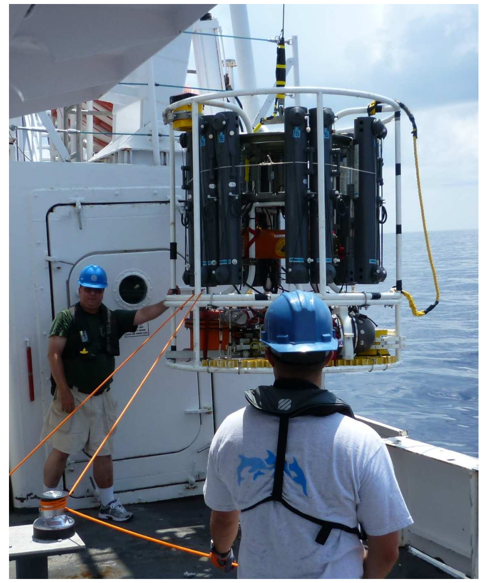

Fig. 2. CTD and rosette being deployed off the R/V *Endeavor*. Non-standard instrumentation including the TETHYS in situ mass spectrometer and an optode affixed to the cage on the bottom of the package. Photo by C. McIntyre, WHOI

both to the WSW of the site (Fig. 4), encountered water enriched with hydrocarbons significantly above background levels. Together, the CTD and *Sentry* resulted in a detailed, multimodal picture of a coherent subsurface plume extending at least 35 km from the *Deepwater Horizon* site. The precise nature of that plume, including constituents, their absolute concentrations, and the fraction of oil confined to the plume are discussed in a forthcoming publication in the scientific literature [6].

A. Acoustic Telemetry

Seawater rapidly attenuates electromagnetic radiation, rendering radio frequency signals of the type commonly employed in terrestrial communications systems ineffective underwater. Acoustic modem systems, e.g. [7], provide a relatively low-bandwidth (40 km $\cdot$ kpbs [8]) alternative that nevertheless allows for the real-time downlink of a portion of data collected subsea and for the transmittal of high-level control commands to the vehicle. *Sentry* uses a commercial acoustic modem integrated into the USBL navigation system.

This system enables commands to be sent to the vehicle and data sent from the vehicle to the surface vessel. The amount of data being obtained on the AUV far exceeds the available bandwidth, so we employ a queuing system in which the user sends a message to the vehicle requesting which information should be transmitted back to the surface vessel for human interpretation. In addition to sending information requests to the vehicle, we can transmit mission re-specification commands to achieve tasks such as changing vehicle depth and retasking the vehicle on new trajectories. This architecture allows us

[2] `http://mw1.google.com/mw-earth-vectordb/disaster/gulf_oil_spill/kml/noaa/nesdis_anomaly_rs2.kml`

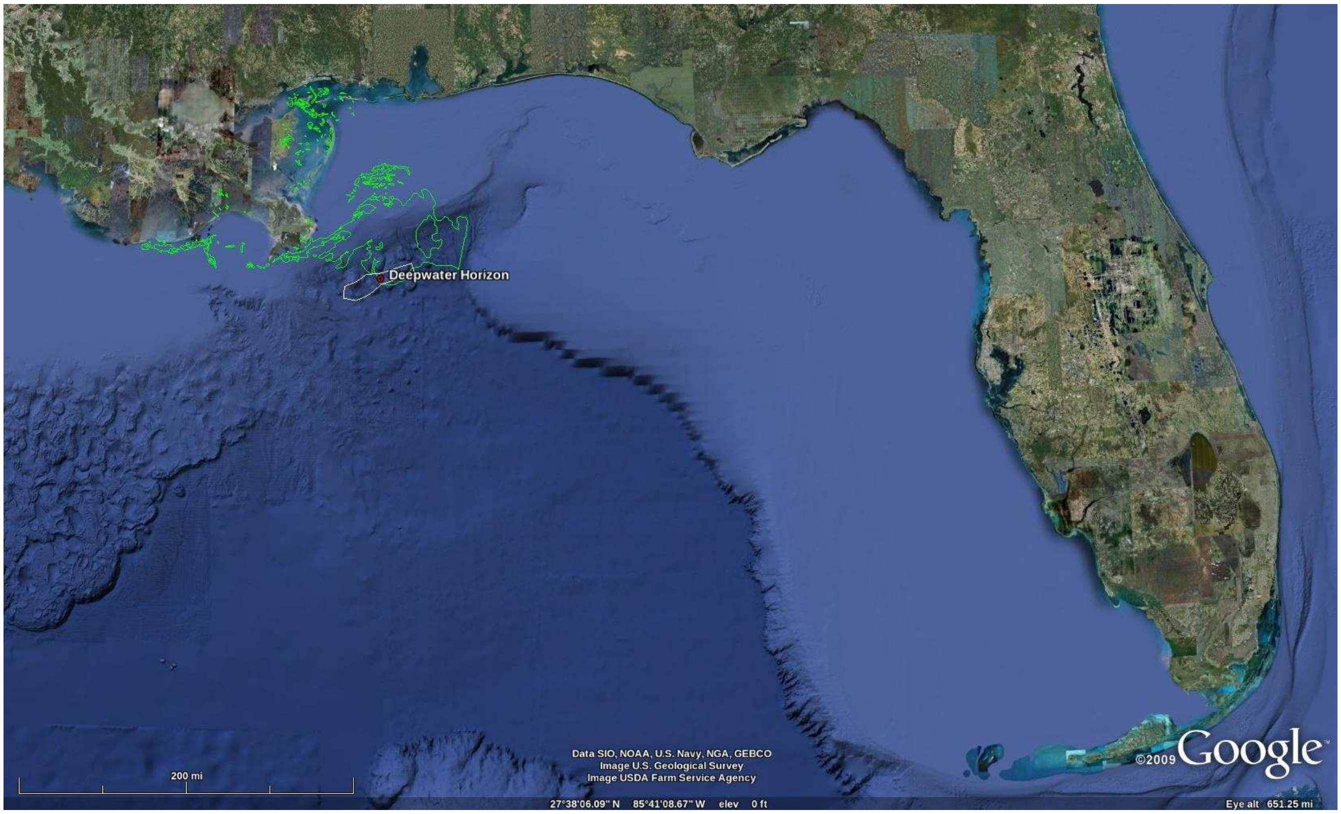

Fig. 3. Work site location in the Gulf of Mexico off the Louisiana coast. All AUV and shipboard operations were carried out within the white bounds indicated, an area about 50 nautical miles in length. For comparison, the light green outlines show the potential oiling footprint of the surface plume observed by NOAA 2010-07-07[2]. The *Deepwater Horizon* site is shown located at $28^o\ 44.3071'$ N, $88^o 21.9611'$ W, based on ship's radar while working near the site.

to receive crucial sensor data and, based on the information obtained, to retask the vehicle in response.

B. Real-time Visualization

Limited cruise duration, limited a priori information about the plume, and the need to obtain precisely targeted water samples from within a dynamic phenomenon required the rapid analysis and visualization of data to devise appropriate sampling strategies and use the available assets efficiently. Leveraging previous work employing the Keyhole Markup Language (KML) for the dissemination and visualization of geo-referenced oceanographic data [9], [10], we provided the science party with near real-time displays of integrated chemical tracer data from all instruments on board the CTD and selected ion peaks from the TETHYS mass spectrometer on board *Sentry*. Fig. 5 shows a screenshot of our data visualization part way through AUV dive sentry064, rendered by Google Earth. The image shows aromatic hydrocarbon fluorimetry from several CTD casts and one tow-yo, real-time normalized methane concentration telemetered acoustically from *Sentry*, and real-time water current profiles generated by the ship's acoustic Doppler current profiler (ADCP). This visualization was instrumental in coordinating the sampling strategies of the CTD and *Sentry*. It aided in site selection and survey design, water sample location selection, real-time survey modification, and provided the first visual confirmation of a coherent subsea plume.

C. Sentry Dives 064 and 065 — Tracking a Subsurface Oil Plume

The first challenge in assessing the extent of the subsurface plume was initially locating it. A CTD tow-yo conducted around the periphery of the *Deepwater Horizon* site registered intense hydrocarbon anomalies at 1100 m depth to the west-southwest of the well head and weaker anomalies at the same depth to the northeast. Based on this and other supporting data we planned a series of AUV surveys aimed at tracking the plume down-current of the well head. The goal of these surveys was to determine the horizontal extent of the plume and provide the necessary reconnaissance for targeted water sampling.

The *Sentry* AUV was deployed on two dives — sentry064 and sentry065 — during which the vehicle tracked the plume over 30 km down-range from the origin of the plume at the well head (Fig. 4). While both dives had identical goals, the manner in which they were conducted differed, and the contrasting survey techniques used in either case each possessed advantages and disadvantages. We planned and executed sentry064 in the conventional manner, with the vehicle following a series of preplanned tracklines designed to repeatedly cross

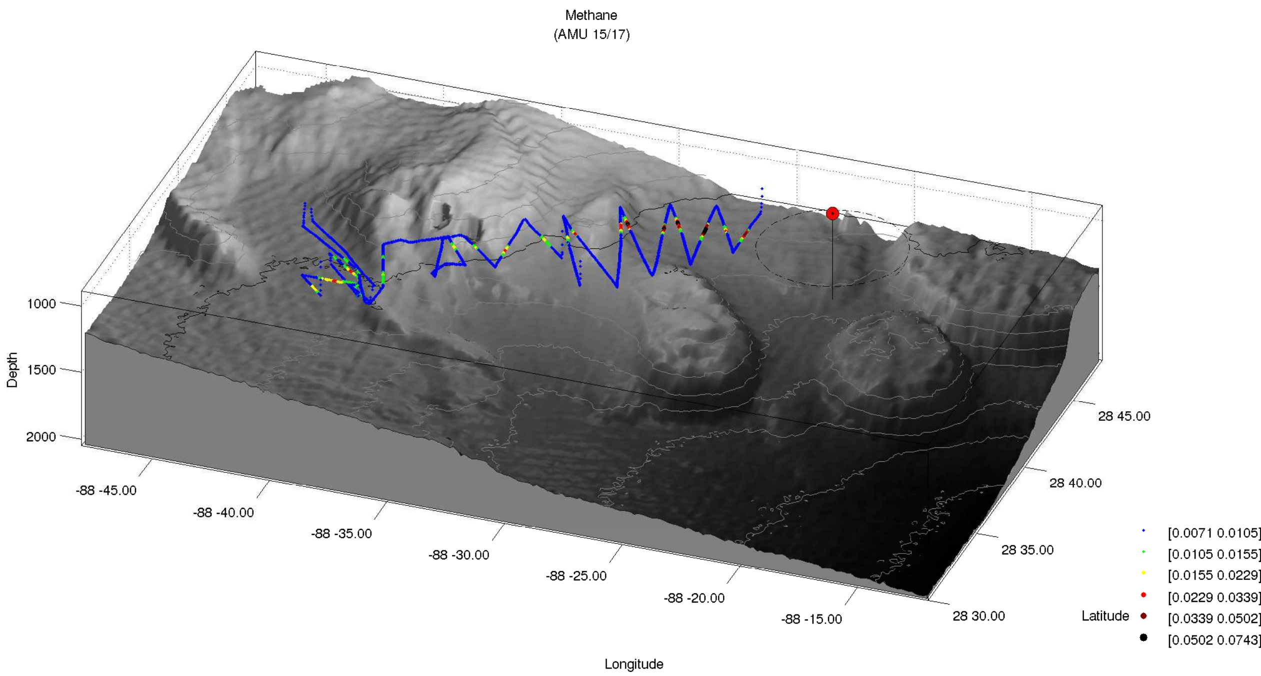

Fig. 4. Normalized Methane observed from the TETHYS mass spectrometer aboard the *Sentry* AUV during two dives, sentry064 and sentry065, to the west of the *Deepwater Horizon* site. The *Deepwater Horizon* site and 5 km exclusion zone are indicated in the perspective view.

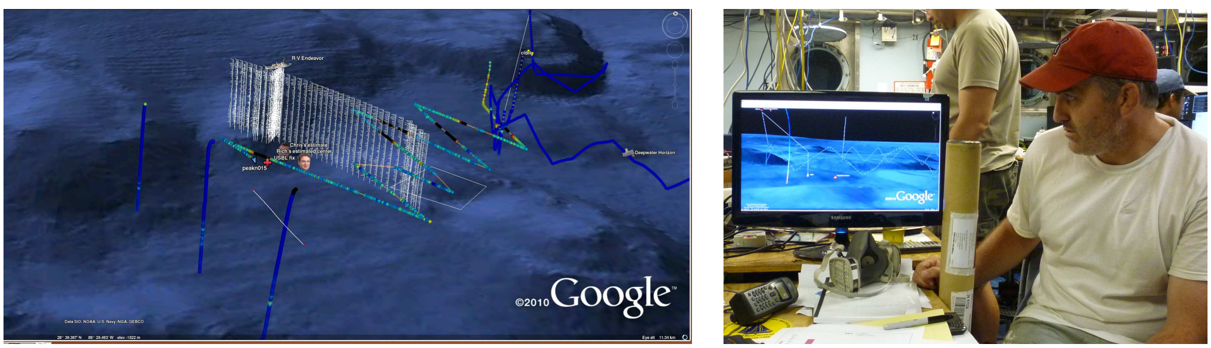

(a) CTD and *Sentry* data rendered in Google Earth.

(b) Co-chief scientist Dr. Reddy considering the data.

Fig. 5. Real-time data visualization: (a) Screenshot of Google Earth rendering taken during the cruise and showing both fluorometer data collected over the preceding days using CTD casts and tow-yos as well as TETHYS mass spectrometer data being telemetered acoustically from the AUV in real time; (b) inspecting the visualisation. Prompt, effective visualization improved the ability of the science party to coordinate sampling and survey activities as well as to alter survey plans in real time, enhancing the efficiency and effectiveness of operations.

the plume at a constant depth. Real-time acoustic telemetry from the vehicle was used to select the site for a CTD cast that was then conducted during the dive and out of acoustic range of the vehicle.

We designed sentry065 similarly but with the intention of acoustically manipulating the mission plan in real time. We planned to cut tracklines short after mass spectrometry data received acoustically indicated a return to background values following a transect of the plume. This strategy was designed to increase the down-current extent of the survey and was employed successfully early during the dive. Plume intensity on later tracklines exhibited an unexpected decrease in magnitude, prompting us to dramatically alter *Sentry's* mission plan, first to reacquire the plume closer to the well head, and then later to refine the survey depth before continuing with (a modified version of) the original survey plan.

At 30 km from the well head the hydrocarbon anomaly remained well above the detection threshold of the TETHYS instrument on *Sentry*; however deteriorating weather conditions prevented further AUV deployments and ultimately forced an end to scientific operations altogether. In total, dives 64 and 65 spanned a total of 61 hours during which *Sentry* spent 47.4 hours deployed and covered over 170 km.

III. Semi-Autonomous Subsea Robotic Survey

Our real-time interactions with *Sentry* yielded scientifically more productive dives but also required us to engage in low-level data processing and trajectory-level mission re-specification. As subsea robots become more sophisticated and the number of robots concurrently in the water increase, the scope for low-level interactions will decrease commensurately. Human oversight will remain valuable but must transition to higher-level interaction. This will require enhanced autonomy on the part of the robots themselves. In this section we discuss the performance of a semi-autonomous method for subsea robotic survey applied, via data-denial simulation, to dive sentry064.

Our method applies the classical sense-plan-adapt (SPA) approach to robotic decision making but with high-level human input at each stage of the cycle. Our primary aim is to reduce the cognitive load on human operators while still leveraging human skill in high-level decision making. This aim aligns well with the reality of limited bandwidth acoustic communications—data pre-processing carried out autonomously subsea can reduce the bandwidth required to telemeter data to the surface; a robot capable of interpreting high level objectives rather than direct trajectory specification will also likely reduce the bandwidth required to transmit control commands. The motivations behind our particular implementation of each stage of the SPA cycle is discussed subsequently.

a) Sense: Various authors have reported on the use of AUVs to trace and/or map both synthetic and naturally occurring turbulent plumes, e.g. [11]–[15]. A necessary component of any of these methods is a mechanism for deciding what sensor readings represent contaminated plume water versus background water. Such mechanisms become more difficult to construct when, as in our case, the signature of the plume within data from the various sensors available was initially unknown.

The sensor suite on board *Sentry* was chosen by scientists based on expert knowledge of the likely chemical constituents of a subsea hydrocarbon plume; nevertheless, significant uncertainty remained concerning the presence, relative concentrations, and manifestation of these constituents in the sensor data. Ultimately the methane measurement produced by the TETHYS instrument provided the most reliable indication of plume presence; however, this knowledge was unavailable prior to human analysis of all sensor data streams.

Our approach to automated plume detection considers all 11 available scalar sensor data streams together as vector-valued data, sorts these into statistically distinct classes, and relies on human interpretation to provide a semantic label for each class as either plume, background, or other. Parametrized statistical models for each class are learned as part of the procedure, meaning the robot can autonomously apply semantic labels to subsequently acquired data.

The model used for classification in this paper is the Bayesian, non-parametric, Variational Dirichlet Process model (VDP) [16]. This model is a mean-field variational approximation of a Dirichlet Process Mixture Model (DPMM) [17], [18]. Important assumptions made in this paper are that observations are distinctly multimodal, can be represented using a Gaussian Mixture Model (GMM), and are independently and identically

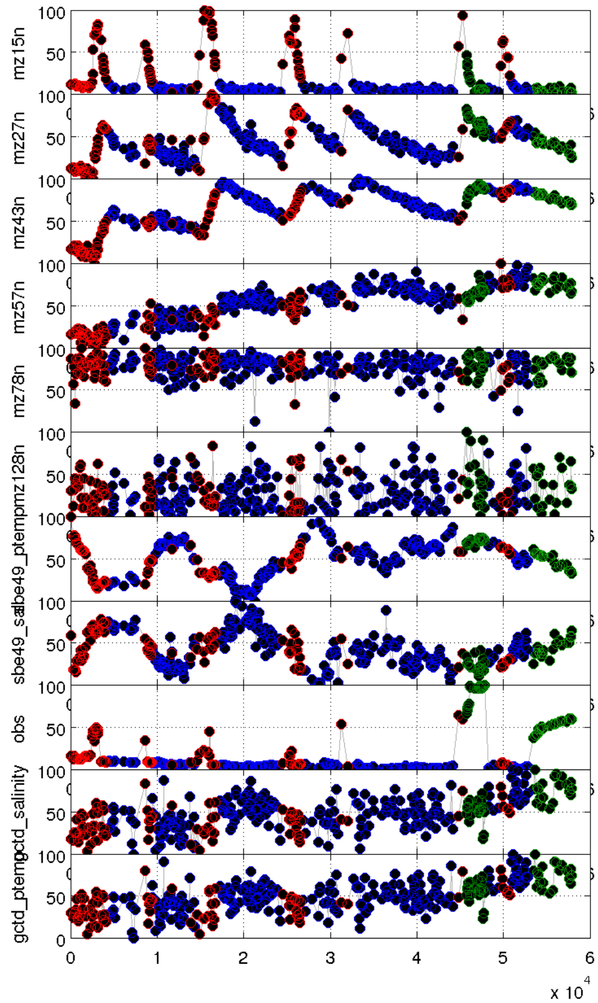

Fig. 6. All 11 scalar chemical sensor data streams interpolated onto the timebase of the TETHYS instrument and classified into four distinct components of a Gaussian Mixture Model. Semantic labeling as three classes, plume (red), background (blue), and other (green, representing two distinct mixture components) was provided by a human.

distributed (i.i.d.) when conditioned on their class label. Its principal feature, besides rapid execution, is that the method automatically infers the number of classes present in the data. Fig. 6 shows the classified output produced after semantic labeling by a human. The algorithm appears to have implicitly identified methane (mz15) and optical backscatter (OBS) as indicative of a distinct class (labeled plume and shown in red), and has also successfully identified two periods of anomalous behavior in the OBS sensor as a distinct classes (labeled other and both shown in green).

As yet our classification process exploits no notion of spatial coherency in the environmental phenomena of interest. While we have attained promising classification results despite this, to adapt vehicle trajectory some mechanism for performing inference over the spatial domain of the survey area is often necessary.

b) Plan: Several spatial inference methods specific to robotic plume mapping as applied especially to plume source localisation exist [14], [19]–[22]. Like [19] our approach employs a Mixture of Gaussian Processes (MGP) to model the spatial coherence of the plume and background; however, because our output space is 11-dimensional rather than a scalar chemical concentration, we perform a logistic regression over the scalar class labels to avoid learning the parameters of what would otherwise become a multivariate MGP. This is known as Gaussian Process Classification (GPC) [23].

Once the hyperparameters of the mixture components have been learned GPC regression provides a way to extrapolate the probability of observing each semantically labeled class to the spatial domain of the survey (Fig. 7). On the basis of this map, an agent can plan by evaluating the expected outcome of future actions relative to a specified objective function, for instance, [24] traded off information gain with traversal cost to generate constrained maximum entropy sampling plans.

In practise, developing good objective functions in the dynamic setting of a scientific expedition remains challenging. On the other hand, scientists and operators try to design pre-programmed AUV surveys in a way that encapsulates key objectives, some of which, like coordination with other assets and weather considerations, would be difficult to encode in a useful objective function because they depend on external circumstances not readily sensed by a deployed robot. We propose that limiting the scope of autonomous planning to modifications of the pre-planned mission can retain good performance relative to these hard-to-codify objectives, and if designed with autonomous adaptation in mind, also benefit from autonomous decision making.

c) Act: Dive sentry064 (Fig 4) was designed to provide multiple down-current horizontal crossings through the hydrocarbon plume, under the assumption that the current would cause the plume to spread along isobaths to the WSW. The increasing amplitude of the zigzag trajectory specified in the mission plan reflected our uncertainty about plume spreading rate and the precise direction of the current. The large amplitude of the survey tracklines in relation to the width of the plume encountered represents an inefficiency that might have been mitigated by terminating tracklines early, as was commanded by human operators via acoustic link on sentry065. This might also have been accomplished autonomously had an appropriate objective function and set of admissible control actions been available.

To test this supposition, we ran a Markov Decision Process (MDP) over the 4 segments of each trackline, allowing the MDP, upon completion of each segment, to select whether to continue on to the next segment or else abandon the rest of the trackline. The MDP was rewarded for completing segments likely, based on the output of the GPC, to encounter plume, and penalised for completing segments likely to pass exclusively through background. This approach is myopic in the sense that the GPC regression is regarded as truth at each iteration of the MDP.

Fig. 7 shows a snapshot of a simulation produced by denying data to the sensing stage of the SPA cycle from trackline segments aborted by the MDP. At the instant shown the simulation has correctly avoided completing the southern end of a pair of tracklines relatively far from the plume. A good GMM for classification was attained after completion of a few tracklines and did not require relearning or relabeling until the anomalous OBS data appeared later in the dive. The GPC appeared less convincing—although the learned hyperparameters varied little, the correlation length scales determined were short relative to the length scale imposed by inter-trackline spacing on the survey. This lack of predictive certainty distant from existing data may be the source of an observed sensitivity in the decision process to the numerical values of the rewards specified in the objective function.

IV. Discussion and Conclusions

Several aspects of this work bear directly the current and future use of robotics for environmental monitoring:

1) The time and spatial scales associated with dynamic features in an environment should drive sampling plan design as well as the selection of appropriate instrumentation, including the use of autonomous platforms. In our case an AUV offered maneuverability and speed advantages over a cable-lowered CTD for reconnoitering the horizontal extent of a subsurface hydrocarbon plume, but was most effective in concert with the CTD.
2) Real-time transmission of data from autonomous platforms can augment the effectiveness of these platforms by enabling operators to deploy other assets before conditions change. We used the real-time AUV data to inform the sampling strategy of the CTD and to target water samples.
3) Some environmental monitoring tasks are characterized by relatively large swaths of uninteresting terrain. In these circumstances, we stand to gain the most from adaptive survey. In our case, we adapted the cruise plan to data as it became available, and, on a finer scale, we also adapted the AUV's trajectory to more effectively sample the feature of interest.
4) Human intervention may increase the scientific yield of robotic surveys, but any increased value must be traded off against the opportunity cost of demanding a human's attention. There is a pressing need to develop autonomous and semi-autonomous data processing and adaptive survey methods that reflect the real challenges of incompletely characterized environmental phenomena, limited processing power and communications bandwidth, and limited endurance.

Based on these observations, we developed a semi-supervised method for adaptive survey that conceivably could have reduced the time spent by the AUV outside the plume during dive sentry064 without requiring intensive operator interaction. The comparison with sentry065 is instructive. On sentry065 intensive human interaction was required to

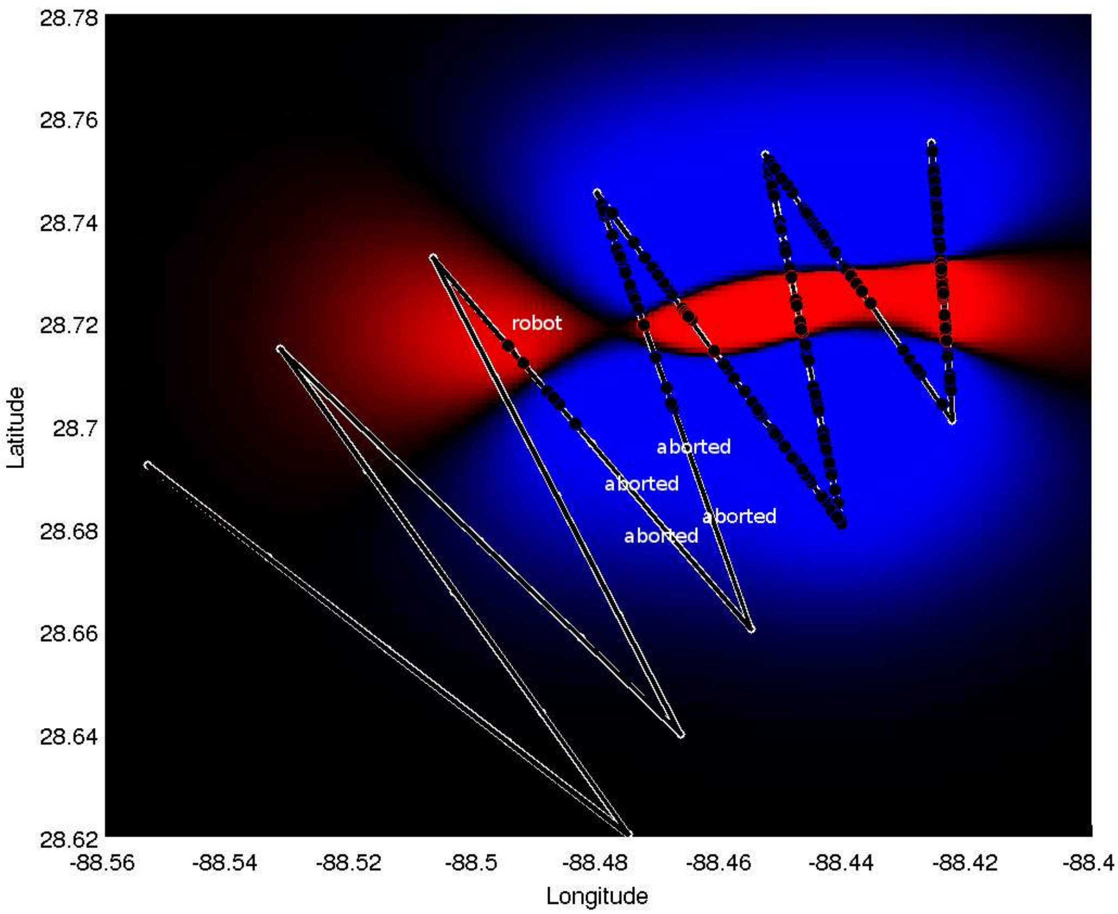

Fig. 7. Data-denial simulation of semi-autonomous adaptive execution of sentry064 part way through the run. The colored circles represent data classified as plume (red) or background (blue). The entire domain of the survey is colored according to the GPC regression, with black representing maximum ambiguity, that is an equal chance of either plume or background.

reacquire the plume signal. Our approach to semi-autonomous adaptive survey relies on a sensible pre-planned mission. Radical changes to the mission plan like that required in sentry065 would require a more complex 3-dimensional environmental model, a far more complete set of admissible control actions, and a consequently much more complex decision process.

ACKNOWLEDGMENT

This work was funded through the National Science Foundation's RAPID program. We acknowledge the captain and crew of the R/V *Endeavor* for their careful handling of *Sentry* during over-the-side operations as well as other members of the scientific party aboard for their management of CTD operations and water sample handling. We also acknowledge the Dorado AUV team from MBARI and the scientific party from the NOAA ship R/V *Thomas Jefferson*, and Prof. Ruoying He of the Dept. of Marine, Earth and Atmospheric Sciences at North Carolina State University, all of whom graciously provided their results from their own investigations into the oil spill. As always, the ABE/*Sentry* team provided excellent vehicle operational support and we are grateful to Andy Billings, Alan Duester, and Scott McCue for their efforts before and during the cruise.

REFERENCES

[1] J. Gillis, "Giant plumes of oil forming under the gulf," New York Times, May 15 2010, http://www.nytimes.com/2010/05/16/us/16oil.html.

[2] MBARI Dorado AUV group, "Data Report: June 2-3, 2010 Operations of MBARI AUV R/V Gordon Gunter Cruise, Gulf of Mexico," Monterey Bay Aquarium Institute, Tech. Rep., June 11 2010.

[3] P. S. Daling, M. O. Moldestad, O. Johansen, A. Lewis, and J. Rodal, "Norwegian testing of emulsion properties at sea–the importance of oil type and release conditions," *Spill Science and Technology Bulletin*, vol. 8, no. 2, pp. 123–136, Apr. 2003.

[4] J. W. Farrington, "NOAA Ship RESEARCHER/Contract Vessel PIERCE Cruise to IXTOC-1 Oil Spill: Overview and Integrative Data Assessment and Interpretation," Woods Hole Oceanographic Institution, Tech. Rep., 1983.

[5] R. Camilli and A. Duryea, "Characterizing spatial and temporal variability of dissolved gases in aquatic environments with in situ mass spectrometry," *Environmental Science and Technology*, vol. 43, no. 13, pp. 5014–5021, 2009.

[6] R. Camilli, C. Reddy, D. Yoerger, B. V. Mooy, M. Jakuba, J. Kinsey, C. McIntyre, S. Sylva, and J. Maloney, "Tracking hydrocarbon plume transport and biodegradation at Deepwater Horizon," Submitted to *Science*.

[7] L. Freitag, M. Grund, S. Singh, J. Partan, P. Koski, and K. Ball, "The WHOI micro-modem: an acoustic communications and navigation system for multiple platforms," in *Proceedings IEEE/MTS Oceans Conference and Exhibition*, vol. 2, 2005, pp. 1086–1092.

[8] J. Partan, J. Kurose, and B. N. Levine, "A survey of practial issues in underwater networks," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 11, no. 4, pp. 23–33, Oct. 2007, special section on ACM WUWNet 2006.

[9] C. Murphy, "Lossy compression and real-time geovisualization for ultra-low bandwidth telemetry from untethered underwater vehicles," Master's thesis, MIT / WHOI Joint Program in Oceanography, Cambridge, MA, Sep. 2008.

[10] C. Murphy and H. Singh, "Human-guided autonomy for acoustically tethered underwater vehicles," in *Proceedings IEEE/MTS Oceans Conference and Exhibition*, Sep. 2008, pp. 1–8.

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

[12] D. Yoerger, A. Bradley, C. G. M.V. Jakuba, T. Shank, and M. Tivey, "Autonomous and remotely operated vehicle technology for hydrothermal vent discovery," *Oceanography*, vol. 20, no. 1, pp. 152–161, Jan. 2007.

[13] M. Fisher and M. Dock, *Trace Chemical Sensing of Explosives*. Wiley-Interscience, 2007, ch. Detection of Trace Explosive Signatures in the Environment, pp. 133–150, ronald L. Woodfin, editor.

[14] M. Jakuba and D. Yoerger, "Autonomous search for hydrothermal vent fields with occupancy grid maps," in *Proceedings Australian Conference on Robotics and Automation*, Dec. 2008.

[15] R. N. Smith, J. Das, Y. Chao, D. A. Caron, B. H. Jones, , and G. S. Sukhatme, "Cooperative multi-auv tracking of phytoplankton blooms based on ocean model predictions," in *Proceedings IEEE/MTS Oceans Conference and Exhibition*, Sydney, Australia, May 2010.

[16] K. Kurihara, M. Welling, and N. Vlassis, "Accelerated variational Dirichlet process mixtures," *Advances in Neural Information Processing Systems*, vol. 19, p. 761, 2007.

[17] T. Ferguson, "A Bayesian Analysis of some Nonparametric Problems," *The Annals of Statistics*, vol. 1, no. 2, pp. 209–230, 1973.

[18] C. E. Antoniak, "Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems," *The Annals of Statistics*, vol. 2, no. 6, pp. 1152–1174, 1974.

[19] C. Stachniss, C. Plagemann, and A. Lilienthal, "Learning gas distribution models using sparse Gaussian process mixtures," *Autonomous Robots*, vol. 26, no. 2, pp. 187–202, 2009.

[20] A. J. Lilienthal and T. Duckett, "Gas source localization by constructing concentration gridmaps with a mobile robot," in *Proceedings ECMR 2003*, 2003.

[21] J. A. Farrell, S. Pang, and W. Li, "Plume mapping via hidden markov methods," *IEEE Transactions on System, Man, and Cybernetics*, vol. 33, no. 6, pp. 850–863, 2003.

[22] S. Pang and J. A. Farrell, "Chemical plume source localization," *IEEE Transactions of Systems, Man, and Cybernetics–Part B: Cybernetics*, vol. 36, no. 5, pp. 1068–1080, 2006.

[23] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Massachusetts Institute of Technology: the MIT Press, 2006.

[24] D. Thompson and D. Wettergreen, "Intelligent maps for autonomous kilometer-scale science survey," in *International Symposium on Artificial Intelligence, Robotics and Automation in Space (ISAIRAS)*, 2008.