Sensor Data Fusion in Marine Robotics*

John J. Leonard and Christopher M. Smith

Department of Ocean Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139, U.S.A.

Abstract - Sensor data fusion is an important and difficult requirement in marine robotics. This paper examines the role of sensor fusion in three different application areas: navigation of autonomous underwater vehicles, acoustic scene reconstruction, and ocean data assimilation. The research issues encountered in these problems include management of uncertainty, modeling of sensor physics, vehicle and environment dynamics, feature extraction, data association, state estimation, and sensor management. Previous research in these areas is reviewed and suggestions for future research are presented.

Keywords:

underwater robotics; data fusion; sonar; navigation.

INTRODUCTION

Improved methods for sensor data interpretation are critical for realization of the full potential of today's state-of-the-art marine robotics technology. Marine robotics technology has undergone a phase of dramatic increase in capability in recent years. This is demonstrated by the recent impressive accomplishments of underwater vehicles such as Jason [2], the Autonomous Benthic Explorer (ABE) [52], Theseus [11], and the autonomous underwater vehicle (AUV) Odyssey II [6]. In concert with this advance in vehicle technology, the capabilities of commercial sensing devices have improved, yielding devices of higher quality with lower weight, cost and power consumption. As these trends continue, the primary challenges in subsea robotics are shifting from development of system and sensor technology to problems of *information processing* and *sensor-based control*. State-of-the-art marine robot systems can produce vast amounts of data, but our ability to use marine robot systems to their full potential is limited by the difficulty of processing this information in autonomous fashion.

This paper examines the role of sensor fusion in three different application areas: navigation of autonomous underwater vehicles, acoustic scene reconstruction, and ocean data assimilation. The next three sections present a description of each of these problems. Subsequently, some of the fundamental research issues shared by these application areas are described and related research in these topics is reviewed. Finally, suggestions are made for future research.

AUV NAVIGATION

Navigation is a critical requirement for any type of mobile robot, but this is especially true for autonomous underwater vehicles. Good navigate information is essential for safe operation and recovery of an AUV, especially for under-ice deployments [4] or in regions of high currents [51]. In addition, often for the data to be gathered by an AUV to be of military, commercial, or scientific value, the location from which the data has been gathered must be precisely known.

Because the ocean is impenetrable to electromagnetic energy except at very low frequencies, navigation systems such as LORAN and GPS are unavailable, and instead acoustic systems must be used [5]. One method for acoustic navigation is long baseline (LBL) navigation, in which the vehicle operates within a pre-calibrated array of acoustic beacons. This technique has been in use for several decades [27], [41], primarily for tethered and manned underwater systems. Because multipath acoustic propagation can result in a high degree of spurious data, reliable operation for autonomous systems remains a challenge.

For an autonomous LBL navigation system, rejection of outliers (data association) is an important function [62]. The AUV must make decisions regarding the quality of the data without the aid of a human operator. The reliability of the solution must be very high, because a survey AUV like the Odyssey II cannot "stop and think" when confronted with a difficult situation; the vehicle must be continually in motion

*To appear in ISOPE-97

to maintain control authority. If the navigation algorithm were to go awry during sea-trials, sending the vehicle hundreds of meters off-course, loss of the vehicle is a very real possibility. Additionally, the navigation algorithm must execute in real-time on a vehicle with limited computational resources.

A significant difficulty in acoustic navigation can be caused by an error in the assumed sound speed profile. Even if the sound speed profile is known at the start of an AUV mission, the acoustic propagation environment can change during the mission. To address this issue to provide navigation in complex and dynamic acoustic environments, Deffenbaugh has developed a technique for very long baseline navigation [17, 18]. The approach uses the extra information provided by the multipath arrivals to invert for sound speed profile variations in space and time, and in the process provide a more accurate position estimate.

For some applications of AUVs, the use of acoustic beacons is undesirable or impractical. If an accurate *a priori* map of the environment is available, one approach to globally-referenced position estimation is to use measurements of geophysical parameters, such as bathymetry, magnetic field, or gravitational anomaly [58, 59, 40, 23]. These approaches are based on matching sensor data with an *a priori* environment map, under the assumption that there is sufficient spatial variation in the parameter(s) being measured to permit accurate localization.

In practice, an up-to-date, high-quality map may be unavailable in the operating area of interest. This motivates research into the problem of concurrent mapping and localization. The goal of concurrent mapping and localization is for the AUV to build a map of its environment and to use that map to navigate in real time. While this has not to our knowledge been attempted with underwater vehicles, the problem has seen some attention in the land robotics community [34]. However, with a few exceptions [44, 49], implementations of concurrent mapping and localization have been restricted to simulation.

A seminal technique for concurrent mapping and localization, called the stochastic map, was published by Smith, Self, and Cheeseman [53]. The stochastic map consists of a single state vector that represents the estimates of the vehicle and feature locations and an associated covariance matrix. As the vehicle moves around its environment, taking measurements of environmental features, the stochastic map is updated using an extended Kalman filter. Moutarlier and Chatila [44] and Rencken [49] have implemented suboptimal versions of the stochastic map using land robots.

The major problems encountered by the stochastic map are the failure of the extended Kalman filter to properly track the highly nonlinear transformations involved in geometric estimation and the fact that the technique scales (at best) quadratically with the number of features present. Also, Smith, Self, and Cheeseman ignored errors which may arise from ambiguity in the source of sensor data. Some of these drawbacks have recently been considered by Uhlmann [61] in a theoretical investigation of concurrent mapping and localization. However, there are many important issues to address in future research to realize this capability on-board an AUV.

ACOUSTIC SCENE RECONSTRUCTION

Sensing the environment is a critical requirement in marine robotics. While a variety of optical sensors are available for underwater use [56, 46], sonar is a natural choice for investigating the turbid marine environment due to its superior propagation characteristics. The aim of acoustic scene reconstruction is to produce a model of an unknown underwater environment based on sonar data obtained from multiple vantage points.

Scene reconstruction is related to the problem of acoustic object recognition, which attempts to identify objects of interest based on known models. Both signal processing and image processing approaches have been pursued. Signal processing techniques attempt to recognize objects by matching templates or resonant signatures of a sonar waveform [16, 30]. Image processing approaches aim to imitate the methods by which skilled human operators interpret visual sonar displays [57].

There are two different approaches to 3-D acoustic scene reconstruction. Stochastic backprojection builds a grid-based, volumetric map of the environment by probabilistically estimating measurement sources [55]. Feature-based reconstruction attempts to associate sonar returns with discrete features in the environment, building a model of the environment composed of geometric primitives [42]. In both approaches, computational complexity is a serious concern.

Stewart pioneered the stochastic backprojection approach and has applied it to a variety of data sets from real ocean settings, such as profile data from the USS *Monitor* [55]. The representation employed is a grid-based, volumetric representation, similar to the occupancy grid developed by Moravec and Elfes for land robot navigation [21]. Singh has extended this work through introduction of a measure of entropy to monitor the amount of information conveyed by new sensor readings in construction of a map [52].

In feature-based scene reconstruction, the first step in interpretation is the extraction of geometric features from raw sensor data, based on a model of sensor physics [35]. After features have been extracted from sensor data, the next step is data association: grouping measurements that share a common origin and rejecting spurious measurements [14]. Once measurements have been partitioned into sets that originate from common environmental features, the shape model (a map of geometric primitives) can be updated [42]. Results to date have been limited to effectively two-dimensional scenes.

OCEAN DATA ASSIMILATION

Ocean data assimilation is the process of integrating real data with a model of an ocean process [38, 50]. One of the largest problems in accurately applying models of ocean flow (or meteorological models, for that matter) is that, because of their dynamic three-dimensional nature and the fact that energy is present at multiple (non-separable) scales, any specification of initial conditions for a forecast will ultimately become uncorrelated with the state of the system [26]. This is due to unmodeled dynamics (such as those at scales which are not modeled for reasons of computation efficiency) and the spread of errors in the initial state. Ocean data assimilation attempts to rectify this inadequacy by providing additional information as the system evolves.

Two different scenarios for ocean data assimilation are *model validation* and *model registration*. In model validation, the objective is to test the validity of a proposed model of an ocean process by comparison with real data. An example of model validation is provided by the proposed study of convective overturning in the Labrador Sea [28, 63]. In model registration, there is a model that is believed to be accurate and the objective is to integrate real data into the model to come up with a specific state estimation for the process of interest. Examples of this are found in numerical weather prediction [26] and oceanographic process field estimation [50].

Ocean data assimilation is a daunting task for two reasons. First, the theoretical and computational challenges of modeling the ocean are tremendous [66]. In addition, our opportunity to observe the ocean is extremely limited. To address the latter issue, the autonomous ocean sampling network (AOSN) project is an ambitious effort that attempts to create a fundamentally new capability to observe the oceans [15]. The idea behind AOSN is to establish a network of small, low cost AUVs supported by acoustic and satellite communication links and power recharging stations. Development of adaptive sampling strategies for real-time ocean data assimilation is an important component of this effort [64].

Munk, Worcester, and Wunsch have pioneered the method of ocean acoustic tomography, which employs acoustic arrays to make precise travel time measurements of a parcel of ocean [45]. In comparison to the point measurements provided by conventional ocean sensing technologies, tomographic measurements are spatially integrating [45]. Recently, the Haro Strait experiment combined moving source tomography, AUV sampling, acoustic propagation modeling, and ocean field prediction [51]. In the future, there promises to be a number of exciting field deployments that will combine the latest technologies for modeling and observing the ocean.

RESEARCH ISSUES

The three applications described above are quite distinct. We believe, however, that these problems have an underlying structural commonality; navigation of AUVs, acoustic scene reconstruction, and ocean data assimilation are all challenging problems in sensor data fusion. The difficulties shared by these problems include the interpretation of uncertain and ambiguous data, modeling of sensor physics, representation of information, and management of computational complexity.

Sensor data fusion can be defined as the purposeful combination of measurements with a priori models, knowledge, and information. Some of the key issues in sensor data fusion are [19, 24, 37]:

- management of uncertainty
- representing, detecting, and tracking features
- modeling of sensor physics
- modeling of platform and environment dynamics
- data association
- state estimation
- sensor management

This section examines these issues from the perspective of marine robotics.

Management of uncertainty

Uncertainty is ubiquitous in marine robotic sensing. In order to properly carry out sensor data fusion, the relative uncertainties of the quantities to be fused must be known or estimated. Uncertainty in marine robotics tasks arises primarily from three sources:

1. uncertainty in the *values* of measurements (noise and biases)
2. uncertainty in the *origins* of measurements (data association error)
3. uncertainty in the *motion* of the sensor or vehicle (navigation error).

The interplay between these different sources of error is what makes sensor data fusion difficult [14]. For example, coupling between navigation and sensing error in concurrent mapping and localization leads to correlations between vehicle and feature position estimates [54, 25, 61].

Since the vehicle must combine uncertain quantities to carry out sensor data fusion, the question of uncertainty representation is an important one. Most data fusion architectures use a Kalman filtering framework for combining measurements and state estimates. This technique is based on a Bayesian representation of uncertainty. Although in a strictly Bayesian paradigm, first-order probabilities are the only meaningful quantities, it is sometimes useful to consider higher-order uncertainty, or the amount of uncertainty in the estimate of the uncertainty of the state estimates [33]. This line of reasoning is followed in the Dempster-Shafer theory of evidence [67] and in other work on higher order probability [13].

In addition to properly combining uncertain data, the vehicle must manage the strategies for addressing these errors so that it can make real-time decisions based on current estimates of its state. Consideration of all three sources of error introduces major computational complexity concerns [14]. Additionally, there may be problems in determining the degree of cross-correlation between uncertain estimates [61].

Representing, detecting, and tracking features

In order to extract useful information from sensor data, the vehicle needs to make some assumptions about its environment. These assumptions form the basis for feature representation. The major dichotomy in feature representation methods is the question of whether to consider the environmental context of the vehicle as one continuous feature (or field), which may be broken down by a grid or pixels, or as a set of discrete features (or objects), which may be independently modeled.

The question of whether a continuous or a discrete geometric representation is more appropriate depends directly on the requirements of the task. For a human-in-the-loop application such as piloting a remotely-operated vehicle (ROV), an occupancy grid offers a method for displaying a large amount of data to the human operator. For object grasping and autonomous recognition of man-made objects, however, an occupancy grid would entail an unduly small cell resolution size, resulting in prohibitive storage and processing requirements. A feature-based geometric approach offers the potential of a compact representation capable of efficiently characterizing a scene and providing direct input to higher level reasoning schemes. This can only be possible, however, if the noise and ambiguity of the input data can be handled in a computationally efficient manner.

Selection of the appropriate features for a given task will depend on a good model of sensor physics as well as environment and platform dynamics. Detection of features based on a sensor model is essentially a problem in data association and track initiation. Tracking features is the subject of state estimation.

Modeling of sensor physics

The key to using any sensor is to have a good sensor model [20]. The objective in developing a sensor model is to support two capabilities:

- Prediction: what data should the sensor produce when observing a known scene from a given position?
- Explanation: given observed sensor data, what is the geometry of the scene that produced the data?

In feature-based acoustic scene reconstruction, a specular sonar model has proven useful in developing geometric constraints for shape estimation. In a specular wavelength regime, isolated features will show up in sonar scans as circular arcs [31]. Circular arcs extracted from sonar data can be used as features for prediction and explanation of sonar data [35].

In acoustic navigation of AUVs, an understanding of acoustic propagation effects is essential, and can be incorporated into the navigation processing to provide robust performance in dynamic environments [17].

Platform and environment dynamics

Dynamic models of the vehicle or sensor platform and the environmental processes being observed are essential in properly representing the dynamic characteristics of the states the vehicle is estimating. Dynamic models also allow the vehicle to make predictions about what it should be sensing, and what is happening in other regions.

One can distinguish between simulation models of vehicle and environment dynamics, which must faithfully reproduce behavior, and models for vehicle use, which must be manageable by the vehicle in real-time. There is the problem of how to most accurately represent these dynamic processes. But there is the further problem of the appropriate model to be used by the AUV; this choice will also depend on computational complexity and speed, fidelity and robustness, operating regime, and expected environmental context.

Providing a model for the vehicle will depend heavily upon the expected range of contexts the vehicle will face. The model should explicitly capture those dynamic phenomena which are of importance or primary interest, while hiding the complexity of the remaining dynamic context.

One example of the integration of environment dynamics with an underwater vehicle simulation occurred in the Haro Strait experiment, in which the AUV Odyssey simulator [7] was coupled with an ocean current model [22] to provide predictions of vehicle performance in high current environments [51].

Data association

Data association is the task of resolving uncertainty in the origins of measurements. The objective is to determine the correspondence between measured data and features of the environment/scene/object being observed, while rejecting spurious measurements. When a continuous feature representation is used, data association is addressed as an integral part of state estimation. For discrete feature models, ambiguity in measurement source must be resolved prior to updating state estimates.

The field of multiple target tracking has provided a number of advanced data association techniques. These can be classified in two categories: target-based and measurement-based. In either approach, the first step is validation gating, a preliminary processing stage in which many infeasible matches between measurements and targets are rejected [3, 61].

An example of a target-based approach is the joint probabilistic data association filter [3], which provides state estimates based on the weighted sum of all the measurements that are in the validation region of a target. These techniques have fixed computational requirements, but produce estimates that average good and bad data. In addition, track initiation for new targets is not explicitly incorporated.

Multiple hypothesis tracking is a measurement-based approach to grouping measurements that originate from the same geometric feature, while rejecting spurious measurements [48]. Each new observation initializes a tree of possible interpretation hypotheses that classify the measurements according to different target models. The tree is grown as new observations are validated with hypothesized targets, and is subsequently pruned to choose the single best interpretation of all past measurements [43, 32]. An important feature of the approach is that ambiguous assignment decisions can be deferred until more data is acquired and a better decision can be made. Cox and Leonard applied multiple hypothesis tracking to land robot map building in a dynamic environment [14]. Moran and Leonard extended the approach to apply to underwater acoustic shape reconstruction of curved objects in two dimensions [42]. Because of the multiple state estimates involved in multiple hypothesis tracking, computation complexity and scaling are significant problems. Current topics of research include pruning strategies and track initialization.

State estimation

Once any necessary data association decisions have been made, state estimation techniques provide the mechanism of integrating sensor measurements obtained from multiple sensors and/or sensing locations. Techniques for robust and efficient state estimation are a central part of research in sensor data fusion [24]. Many approaches are based on the Kalman filtering [3] and Bayesian decision theory [9].

Two important issues are the inherently nonlinear relationships between the parameters being estimated, and correlations between state estimates. Correlations may be due to *a priori* assumptions, correlated measurements, and correlated noise [65]. These correlations may be explicitly modeled, but it should be pointed out that some correlations may be unobservable (and hence unmodelable). For situations when unmodeled correlations are present, Uhlmann has developed an alternative technique for estimation updating called Covariance Intersection [61]. This technique has important implications for concurrent localization and mapping.

Sensor management

Sensor management is the problem of deciding what to sense and from where to sense. Manyika writes "when presented with several sensing options or configurations, the option making the best use of sensor resources to achieve sensing goals must be chosen. [39]" For example, an underwater robot trying to build a map of an unknown environment needs to decide where to move in order to best improve the map it is creating.

In some cases, for example when a simple uniform survey is adequate, sensor management can be performed prior to the mission. However, when the environment is changing or discrete features are being sought (e.g., a search and mapping mission), an adaptive approach to sensor management is needed.

Adaptive sensor management has been studied under many guises, such as action selection [60], active perception [1], and adaptive sampling [8, 64]. A general formulation of adaptive sensor management is hard to obtain because the metrics by which one would evaluate different sensing strategies are highly dependent on the sensor employed and the task at hand.

Manyika has investigated sensor management in the context of decentralized sensor fusion architectures [39]. This work relates to marine data fusion in the context of using multiple AUVs for cooperative mapping. The use of multiple AUVs raises the issue of communication management, because of the bandwidth limitations, time delays, and unreliability of underwater communication [12, 29].

FUTURE RESEARCH

We believe that the problem of enabling an AUV to navigation in, and build a map of, an unknown environment provides a great challenge which can focus future research efforts. Realization of this capability — concurrent mapping and localization — would draw on recent advances in all the research issues discussed above. The reliable extraction of feature location estimates from sonar data in natural environments presents a formidable challenge. However, experiments can begin in environments with man-made objects or distinctive landmarks.

Data association remains a key source of computational complexity. To this point, *ad hoc* approaches to data association have not proven to be robust. It may be possible that multiple hypothesis tracking can offer a foundation on which to develop a comprehensive framework for management of uncertainty in robotic sensing. Attaining this goal will require substantial innovations in the approach, incorporating new methods for non-linear representation and track initiation. A fundamental issue is the initialization of 3D geometric primitives when individual measurements are sparse and provide weak geometric constraints. Because multiple hypothesis tracking is exponentially complex, improved computational efficiency will be essential.

Advanced simulation environments can be a useful aid in developing new methods for data fusion. Simulation has been vital to developers of AUVs because it offers a cost effective alternative to expensive and hazardous field testing [47, 10]. A virtual ocean environment goes beyond the simulation of the AUV itself to incorporate simulations of sensor physics and ocean environment dynamics, with realistic error models. While simulation is no substitute for real data from a deployment at-sea, components of a virtual environment can serve as a human-computer interface for a real experiment [36].

SUMMARY

This paper has surveyed some of the problems of sensor data fusion in marine robotics. Advances in this field have the potential to play a key role in marine applications such as climate change assessment, marine habitat monitoring, oil exploration, and underwater inspection and repair. The high costs and complexity of ocean operation result in a high potential payoff when robotic technology can be successfully applied.

ACKNOWLEDGEMENTS

This work was supported in part by the Naval Undersea Warfare Center, Newport, Rhode Island, under contract N66604-94-C-D057.

REFERENCES

[1] R.K. Bajcsy. Active perception. *Proceedings of the IEEE*, 76(8):996–1005, August 1988.

[2] R. Ballard. Soundings: The Newsletter of the JASON Project for Education, Waltham, MA, http://seawifs.gsfc.nasa.gov/scripts/JASON.html, 1996.

[3] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association.* Academic Press, 1988.

[4] J. G. Bellingham, C. Chryssostomidis, M. Deffenbaugh, J. J. Leonard, and H. Schmidt. Arctic under-ice survey operations. In *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, pages 50–59, 1993.

[5] J. G. Bellingham, T. R. Consi, U. Tedrow, and D. Di Massa. Hyperbolic acoustic navigation for underwater vehicles: Implementation and demonstration. In *Proceedings AUV '92*, pages 304–309, 1992.

[6] J. G. Bellingham, C. A. Goudey, T. R. Consi, J. W. Bales, D. K. Atwood, J. J. Leonard, and C. Chryssostomidis. A second generation survey AUV. In *IEEE Conference on Autonomous Underwater Vehicles*, Cambridge, MA, 1994.

[7] J. G. Bellingham and J. J. Leonard. Task configuration with layered control. In *Proceedings of the IARP 2nd Workshop on Mobile Robots for Subsea Environments*, pages 193–302, Monterey, CA, USA, May 1994.

[8] J. G. Bellingham and J. S. Willcox. Optimizing AUV oceanographic surveys. In *IEEE Conference on Autonomous Underwater Vehicles*, Monterey, CA, 1996.

[9] James O. Berger. *Statistical Decision Theory and Bayesian Analysis.* Springer-Verlag, 1980.

[10] D. P. Brutzman. NPS AUV integrated simulator. Master's thesis, Naval Postgraduate School, Cambridge, MA, March 1992.

[11] B. Butler. Field trials of the Theseus AUV. In *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, pages 6–15, September 1995.

[12] J. Catipovic. Performance limitations in underwater acoustic telemetry. *IEEE J. Ocean Engineering*, 15(3):205–216, July 1990.

[13] T. Chavez. Modeling and measuring the effects of vagueness in decision models. *IEEE Trans. Systems, Man, and Cybernetics*, 26(3):311–323, May 1996.

[14] I. J. Cox and J. J. Leonard. Modeling a dynamic environment using a Bayesian multiple hypothesis approach. *Artificial Intelligence*, 66(2):311–344, April 1994.

[15] T. Curtin, J. G. Bellingham, J. Catipovic, and D. Webb. Autonomous ocean sampling networks. *Oceanography*, 6(3):86–94, 1993.

[16] C. Dean and M. Werby. Target shape and material composition from resonance echoes of submerged elongated elastic targets. In *Proc. SPIE Conference on Automatic Object Recognition*, volume 1471, pages 54–65, Orlando, FL, 1991.

[17] M. Deffenbaugh. A matched field processing approach to long range acoustic navigation. Master's thesis, Massachusetts Institute of Technology, 1994.

[18] M. Deffenbaugh, H. Schmidt, and J. Bellingham. Acoustic positioning in a fading multipath environment. In *IEEE Oceans*, pages 596–599, 1996.

[19] H. F. Durrant-Whyte. *Integration, coordination and control of multi-sensor robot systems.* Boston: Kluwer Academic Publishers, 1987.

[20] H. F. Durrant-Whyte. Sensor models and multi-sensor integration. *Int. J. Robotics Research*, 7(6):97–113, 1988.

[21] A. Elfes. Sonar-based real-world mapping and navigation. *IEEE Journal of Robotics and Automation*, RA-3(3):249–265, June 1987.

[22] M.G.G. Foreman, R.A. Walters, and R.F. Henry. A model for simulating currents in the eastern juan de fuca and the southern strait of georgia. In *IEEE Oceans*, pages 335–340, 1993.

[23] E. Geyer, P. Creamer, J. D'Appolito, and R. Gains. Characteristics and capabilities of navigation systems for unmanned untethered submersibles. In *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, pages 320–347, 1987.

[24] G. Hager. *Task-directed Sensor Fusion and Planning: A Computational Approach.* Boston: Kluwer Academic Publishers, 1990.

[25] P. Hebert, S. Betge-Brezetz, and R. Chatila. Decoupling odometry and exteroceptive perception in building a global world map of a mobile robot: The use of local maps. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 757–764, April 1996.

[26] J. R. Holton. *An Introduction to Dynamic Meteorology.* San Diego, CA: Academic Press, 1992.

[27] M. Hunt, W. Marquet, D. Moller, K. Peal, W. Smith, and R. Spindel. An acoustic navigation system. Technical Report WHOI-74-6, Woods Hole Oceanographic Institution, 1974.

[28] H. Jones and J. Marshall. Convection with rotation in a neutral ocean: a study of open ocean deep convection. *Journal of Physical Oceanography*, 23:1009–1039, June 1993.

[29] J. H. Kim. Remote monitoring and control of autonomous underwater vehicles. Master's thesis, MIT, 1997.

[30] R. Kuc. Fusing binaural sonar information for object recognition. In *IEEE/SICE/RSJ International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pages 727–735, 1996.

[31] R. Kuc and M. W. Siegel. Physically based simulation model for acoustic sensor robot navigation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9(6):766–778, November 1987.

[32] T. Kurien. Issues in the design of practical multitarget tracking algorithms. In Y. Bar-Shalom, editor, *Multitarget-Multisensor Tracking: Advanced Applications*, pages 43–83. Boston: Artech House, 1990.

[33] P. Lehner, K. Blackmond Laskey, and D. Dubois. An introduction to issues in higher order uncertainty. *IEEE Trans. Systems, Man, and Cybernetics*, 26(3):289–293, May 1996.

[34] J. J. Leonard and H. F. Durrant-Whyte. Simultaneous map building and localization for an autonomous mobile robot. In *Proc. IEEE Int. Workshop on Intelligent Robots and Systems*, Osaka, Japan, 1991.

[35] J. J. Leonard and H. F. Durrant-Whyte. *Directed Sonar Sensing for Mobile Robot Navigation.* Boston: Kluwer Academic Publishers, 1992.

[36] J. J. Leonard, S. T. Tuohy, J. G. Bellingham, B. A. Moran, J. H. Kim, H. Schmidt, N. M. Patrikalakis, and C. Chryssostomidis. Virtual environments for AUV development and ocean exploration. In *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, pages 436–443, September 1995.

[37] Zhi-Quan Luo and John N. Tsitsiklis. Data fusion with minimal communication. Technical Report LIDS-P-2065, MIT Laboratory for Information and Decision Systems, 1991.

[38] P. Malanotte-Rizzoli and A. R. Robinson, editors. *Ocean processes in climate dynamics: global and mediterranean examples.* Kluwer Academic Publishers, 1994.

[39] J. S. Manyika and H. F. Durrant-Whyte. *Data Fusion and Sensor Management: A decentralized information-theoretic approach.* New York: Ellis Horwood, 1994.

[40] M. B. May. Gravity navigation. In *Record of the 1978 Position Location and Navigation Symposium*, pages 212–218, San Diego, CA, USA, November 1978.

[41] P. H. Milne. *Underwater Acoustic Positioning Systems.* London: E. F. N. Spon, 1983.

[42] B. A. Moran, J. J. Leonard, and C. Chryssostomidis. Curved shape reconstruction using multiple hypothesis tracking. *IEEE J. Ocean Engineering*, 1996. To Appear.

[43] S. Mori, C. Chong, E. Tse, and R. Wishner. Tracking and classifying multiple targets without a priori identification. *IEEE Transactions on Automatic Control*, AC-31(5), May 1986.

[44] P. Moutarlier and R. Chatila. An experimental system for incremental environment modeling by an autonomous mobile robot. In *1st International Symposium on Experimental Robotics*, Montreal, June 1989.

[45] W. Munk, P. Worcester, and C. Wunsch. *Ocean Acoustic Tomography.* Cambridge University Press, 1995.

[46] S. Nagahdaripour and S. Zhang. Determining the size and position of cylindrical objects from optical images in underwater environments. In *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, pages 232–242, September 1995.

[47] G. Pappas, W. Shotts, M. O'Brien, and W. Wyman. The DARPA/NAVY unmanned undersea vehicle program. *Unmanned Systems*, 9(2):24–30, 1991.

[48] D. B. Reid. An algorithm for tracking multiple targets. *IEEE Transactions on Automatic Control*, AC-24(6), December 1979.

[49] W. D. Rencken. Concurrent localisation and map building for mobile robots using ultrasonic sensors. In *Proc. IEEE Int. Workshop on Intelligent Robots and Systems*, pages 2192–2197, Yokohama, Japan, 1993.

[50] A. R. Robinson. Physical processes, field estimation, and an approach to interdisciplinary ocean modeling. Technical report, Department of Earth and Planetary Science, Harvard University, October 1994.

[51] H. Schmidt, J. Bellingham, M. Johnson, D. Herold, D. Farmer, and R. Pawlowcisz. Real-time frontal mapping with AUVs in a coastal environment. In *IEEE Oceans*, pages 1094–1098, 1996.

[52] H. Singh, D. Yoerger, R. Bachmayer, A. Bradley, and W. Stewart. Sonar mapping with the autonomous benthic explorer (ABE). In *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, pages 367–375, September 1995.

[53] R. Smith and P. Cheeseman. On the representation and estimation of spatial uncertainty. *International Journal of Robotics Research*, 5(4):56, 1987.

[54] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I. Cox and G. Wilfong, editors, *Autonomous Robot Vehicles.* Springer-Verlag, 1990.

[55] W. K. Stewart. *Multisensor Modeling Underwater with Uncertain Information.* PhD thesis, Massachusetts Institute of Technology, 1988.

[56] W. K. Stewart. Remote sensing issues for intelligent underwater systems. In *International Conference on Computer Vision*, pages 230–235, 1991.

[57] G. Trimble, J. Vilaro, D. Okamura, R. Lum, and K. Dutta. Underwater object recognition and automatic positioning to support dynamic positioning. In *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, pages 273–279, 1991.

[58] S. T. Tuohy, J. J. Leonard, J. G. Bellingham, N. M. Patrikalakis, and C. Chryssostomidis. Map based navigation for autonomous underwater vehicles. *International Journal of Offshore and Polar Engineering*, 6(1):9–18, March 1996.

[59] C. Tyren. Magnetic anomalies as a reference for ground-speed and map-matching navigation. *The Journal of Navigation*, 35(2):242–254, May 1982.

[60] Toby Tyrrell. Defining the action selection problem. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pages 1152–1157, 1992.

[61] J. Uhlmann. *Dynamic Map Building and Localization: New Theoretical Foundations.* PhD thesis, University of Oxford, 1995.

[62] J. Vaganay and V. Rigaud. Supervised navigation: optimal algorithm example and needs for autonomous supervision. In *Proc. IEEE Int. Conf. Robotics and Automation*, May 1995.

[63] M. Visbeck. Labrador sea experiment home page, http://www.ldeo.columbia.edu/~visbeck/labsea/, 1997.

[64] J. Scott Willcox, Yanwu Zhang, James G. Bellingham, and John Marshall. Auv survey design applied to oceanic deep convection. In *IEEE Oceans '96*, 1996.

[65] Alan S. Willsky, Martin G. Bello, David A. Castanon, Bernard C. Levy, and George C. Verghese. Combining and updating of local estimates and regional maps along sets of one-dimensional tracks. *IEEE Transactions on Automatic Control*, AC-27(4):799–813, August 1982.

[66] C. Wunsch. *The Ocean Circulation Inverse Problem.* Cambridge University Press, 1996.

[67] Ronald R. Yager, Mario Fedrizzi, and Janusz Kacprzyk, editors. *Advances in the Dempster-Shafer Theory of Evidence.* John Wiley and Sons, Inc., 1994.