

Feature Extraction and Data Association for AUV Concurrent Mapping and Localisation

I. Tena Ruiz, Y. Petillot, D. M. Lane, C. Salson

Ocean Systems Laboratory

Department of Computing and Electrical Engineering

Heriot-Watt University, Edinburgh, EH14 4AS

Scotland, UK

<http://www.cee.hw.ac.uk/oceans/>

email:(tena,ceeyrp,dml)@cee.hw.ac.uk

Abstract

This paper describes a *Concurrent Mapping and Localisation (CML)* algorithm suitable for localising an *Autonomous Underwater Vehicle (AUV)*. The proposed CML algorithm uses a standard off-the-shelf sonar for sensing the environment. The returns from the sonar are used to detect targets in the vehicle's vicinity. These targets are used in conjunction with a vehicle model by the CML algorithm to concurrently build an absolute map of the environment and localise the vehicle in absolute coordinates. In order for the algorithm to work, the stored targets must be associated to the sonar returns at each iteration. Given the nature of sonar data, false returns complicate this process. The choice of targets and a suitable *data association* strategy is, therefore, vital. The chosen targets consist of returns of a significant strength. The segmentation detects these targets and calculates (a) the relative position of their center of mass with respect to the vehicle, (b) the targets' surface size, and (c) the targets' first invariant moment. This information is used by the system to perform the data association. We have chosen to adapt the well known *Multiple Hypothesis Tracking Filter (MHTF)* [1] to the CML structure. This is a *measurement-oriented* approach that finds the probability that an established target gave rise to a certain return. The paper presents results with real sonar data.

1 Introduction

The advent of AUVs has posed a number of new challenges to the robotics research community. Amongst these challenges the question of true autonomy remains unresolved. Autonomy can be defined as the ability to provide for oneself without the help of others. *Unmanned Underwater Vehicles (UUVs)*, be they AUVs or ROVs, are not yet capable of navigating without external assistance. Navigation itself poses three distinct questions: "where am I?", "where am

I going?" and "how should I get there?" [2]. This paper focuses on the first of these questions.

Most UUVs are equipped with *dead-reckoning sensors*, such as a Doppler Velocity Log (DVL), inertial rate gyros, etc. These types of sensors suffer from drift and the error in the vehicle's position will thus grow without bounds. To fix the position of the vehicle on the world frame *absolute-positioning* sensors are used. Commercial absolute-positioning sensors adapted to the underwater environment include *acoustic positioning* systems; *acoustic super short*, *short* and *long baseline navigation*. All of these systems require the vehicle to be within a volume of water that they cover, therefore restricting the vehicle's exploratory capabilities and not allowing for true autonomy. This has motivated research on CML. The basis of CML is to build a map and concurrently localise the vehicle in the map that is being built [2, 3, 4]. Our chosen approach is to implement the *stochastic map* proposed by Smith, Self and Cheeseman [5]. The stochastic map is essentially an augmented *Extended Kalman Filter (EKF)*. Due to the nature of the algorithm, as the number of targets increase they are simply added onto the state vector and the matrix operations eventually become too costly for real-time implementation. Different methods have been proposed to tackle this particular problem. Of these, covariance intersection [6], decoupled stochastic mapping [7] and the geometric projection filter [8] show potential. Other prominent problems related to the stochastic map approach are the feature extraction and data association. Our previous research on sonar data [9, 10, 11, 12] presented different approaches for tracking returns, for the purpose of classification and obstacle avoidance. This research allowed us to develop a number of algorithms for segmenting sonar returns and tracking identified targets. Our latest research efforts consist in blending this knowledge with CML. The proposed system identifies targets of significant signal strength and feeds the data as-




Figure 1: Overview of the System

sociation algorithm, an adapted version of the MHTF developed by Reid [1]. Targets that have been appropriately associated will be used to update the stochastic map, new targets will be integrated into the stochastic map and returns from unreliable targets will be ignored.

The following section outlines the operation of the different modules of the system and the interactions between these. Section 3 will describe the theoretical principles used by the underlying algorithms. Section 4 describes our implementation. Section 5 will illustrate the results. Finally, section 6 will provide a brief conclusion of our findings and outline our future work.

2 Overview

The main modules of the system can be seen in figure 1. The sonar data is initially treated as an image and processed by means of standard image processing techniques. It is subsequently segmented and the targets are identified. The segmentation process provides the data association, MHTF, with a set of observed relative distances between the targets and the vehicle, along with the observed targets' size and first invariant moment descriptors [13].

The MHTF data association algorithm associates the segmented targets to existing targets on the stochastic map. Segmented targets that are not associated to any of the tracked targets are used to initialise new map targets. False alarms identified in the data association process are discarded.

The stochastic map is updated and maintained in absolute coordinates.

3 Theoretical Principles

In this section we examine the theoretical foundations behind the stochastic map and the MHTF, sections 3.1 and 3.2. This theory forms the backbone of our approach, described in section 4.

3.1 The Stochastic Map

The stochastic map is an augmented state EKF [14, 15]. In this incarnation the filter now holds the relevant states of the vehicle and those of the targets in a single state vector. The advantage of this method is that it allows to continually update and maintain the vehicle-to-vehicle, target-to-vehicle and target-to-target correlations. Recent research [16] has demonstrated a number of benefits obtained by maintaining these correlations, namely:

- the determinant of any submatrix of the map covariance matrix decreases monotonically as observations are made successively,
- in the limit, as the number of observations increases, the target estimates become fully correlated, and
- in the limit, the covariance associated with any single target location estimate is determined only by the initial covariance in the vehicle location estimate.

These benefits have motivated our choice of our approach. Under this architecture the new state vector $\mathbf{x}(\cdot)$ assumes the following form,

$$\mathbf{x}(k) = [\mathbf{x}_v(k) \quad \mathbf{x}_1(k) \quad \mathbf{x}_2(k) \quad \dots \quad \mathbf{x}_n(k)]^T \quad (1)$$

where $\mathbf{x}_v(k)$ holds the state of the vehicle and $\mathbf{x}_1(k), \mathbf{x}_2(k), \dots, \mathbf{x}_n(k)$ hold the states of the n targets. The estimated error covariance for this system,

$$\mathbf{P}(k) = \begin{bmatrix} \mathbf{P}_{vv}(k) & \mathbf{P}_{v1}(k) & \mathbf{P}_{v2}(k) & \dots & \mathbf{P}_{vn}(k) \\ \mathbf{P}_{1v}(k) & \mathbf{P}_{11}(k) & \mathbf{P}_{12}(k) & \dots & \mathbf{P}_{1n}(k) \\ \mathbf{P}_{2v}(k) & \mathbf{P}_{21}(k) & \mathbf{P}_{22}(k) & \dots & \mathbf{P}_{2n}(k) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{P}_{nv}(k) & \mathbf{P}_{n1}(k) & \mathbf{P}_{n2}(k) & \dots & \mathbf{P}_{nn}(k) \end{bmatrix} \quad (2)$$

where the sub-matrices $\mathbf{P}_{vv}(k)$, $\mathbf{P}_{vi}(k)$ and $\mathbf{P}_{ii}(k)$ are the vehicle-to-vehicle, vehicle-to-target and target-to-target covariances respectively.

The state and covariance are updated according to the EKF update equations. The stochastic map assumes fixed targets and the resulting state propagation will become,

$$\hat{\mathbf{x}}_v(k) = \mathbf{f}_v[\hat{\mathbf{x}}_v(k-1), \mathbf{u}(k), \mathbf{0}, k] \quad (3)$$

where $\hat{\mathbf{x}}_v(\cdot)$ is the vehicle's state and $\mathbf{f}_v[\hat{\mathbf{x}}_v(\cdot), \mathbf{u}(k), \mathbf{0}, k]$ is the vehicle's dynamic model. And its associated covariance will be propagated thus,

$$\mathbf{P}(k) = \mathbf{F}_{x_v} \mathbf{P}(k-1) \mathbf{F}_{x_v}^T + \mathbf{F}_{w_v} \mathbf{Q}(k) \mathbf{F}_{w_v}^T \quad (4)$$

where \mathbf{F}_{x_v} is the Jacobian of the vehicle model with respect to the vehicle state, used to linearise the state of the vehicle error $\hat{\mathbf{x}}_v(k-1)$, and \mathbf{F}_{w_v} is the Jacobian of the vehicle model with respect to the process noise.

The prediction for observed target i is written as

$$\mathbf{z}_i(k) = \mathbf{h}_i[\hat{\mathbf{x}}(k), \mathbf{0}, k] \quad (5)$$

where $\mathbf{h}_i[\hat{\mathbf{x}}(k), \mathbf{0}, k]$ is the observation model. Thus the innovation v_i is defined as

$$v_i = \hat{\mathbf{z}}_i(k) - \mathbf{z}_i(k) \quad (6)$$

with innovation covariance, $\mathbf{S}_i(k)$, defined as

$$\mathbf{S}_i(k) = \mathbf{H}_i(k)\mathbf{P}(k)\mathbf{H}_i^T(k) + \mathbf{R}(k) \quad (7)$$

where $\mathbf{H}(k)$ is a matrix holding the Jacobian of the observed target evaluated at the latest estimate of the state and $\mathbf{R}(k)$ is a matrix of the measurement error covariance.

The gain of the filter, \mathbf{K}_i can now be written as,

$$\mathbf{K}_i(k) = \mathbf{P}(k)\mathbf{H}_i^T(k)\mathbf{S}_i^{-1}(k) \quad (8)$$

And the corrected state estimate becomes,

$$\hat{\mathbf{x}}(k+1) = \hat{\mathbf{x}}(k) + \mathbf{K}_i(k)v_i(k) \quad (9)$$

and its associated covariance is updated according to,

$$\mathbf{P}(k+1) = \mathbf{P}(k) - \mathbf{K}_i(k)\mathbf{S}_i(k)\mathbf{K}_i^T(k) \quad (10)$$

3.2 MHTF

The MHTF algorithm calculates the probability that each established target, or a new target, gave rise to a certain observation. The filter works by evaluating hypotheses of all the possible associations, $\Omega(k)$, up to time k .

Hypotheses are made by associating to $\Omega(k-1)$ each observation, $\mathbf{z}_i(k)$. For each observation, $\mathbf{z}_i(k)$, there are three possible associations:

- it belongs to an existing target,
- it is a new target or
- it is a false alarm.

A hypothesis matrix is built where all the possible configurations are considered.

The joint cumulative event, $\Theta(k)_l$, at time k is made up of the joint event, $\Theta(k-1)_s$, and the current association event, $\theta(k)$. The conditional probability of a cumulative event at time k can be written as,

$$P\{\Theta(k)_l|Z(k)\} = P\{\theta(k), \Theta(k-1)_s|Z(k), Z(k-1)\} \quad (11)$$

From this a recursive relationship may be written,

$$\begin{aligned} P\{\Theta(k)_l|Z(k)\} = \\ \frac{1}{c} p[Z(k)|\theta(k), \Theta(k-1)_s|Z(k-1)] \\ P\{\theta(k)|\Theta(k-1)_s, Z(k-1)\} P\{\Theta(k-1)_s|Z(k-1)\} \end{aligned} \quad (12)$$

This can be shown to be [15],

$$\begin{aligned} P\{\Theta(k)_l|Z(k)\} = \\ \frac{1}{c} \frac{\phi! \nu!}{m_k} \mu_F(\phi) \mu_N(\nu) V^{-\phi-\nu} \prod_{i=1}^{m_k} [N_{\tau_i}[\mathbf{z}_i(k)]]^{\tau_i} \quad (13) \\ \prod_t (P_D^t)^{\delta_t} (1 - P_D^t)^{1-\delta_t} P\{\Theta(k-1)_s|Z(k-1)\} \end{aligned}$$

where ϕ , τ and m_k are respectively the number of false alarms, new targets and measurements in the event $\theta(k)$, $\mu_F(\phi)$ and $\mu_N(\nu)$ are the densities of false and new targets respectively, V is the hypervolume of the surveillance region, N signifies the normal law and τ_i is an indicator variable of value one if measurement $\mathbf{z}_i(k)$ came from an established track and zero otherwise, t is the number of targets, P_D^t is the probability of detecting prior targets and δ_T is an indicator variable which is of value one if target t is detected at time k and zero otherwise.

The conditional probabilities of each cumulative event are calculated and the most probable event is selected as the valid hypothesis.

4 Implementation

We now describe the proposed system. Section 4.1 relates certain aspects relevant to the pre-processing of sonar data. In section 4.2 we describe the processing and segmentation of the sonar returns. A description of the changes made to the MHTF algorithm to fit into our structure is given in section 4.3. Finally particular aspects of our stochastic map algorithm are examined in section 4.4.

4.1 Sonar Data

Each scanned sector returned by the sonar is treated as an image by the segmentation algorithm. This is a reliable process as most electronically scanned sensors allow for update rates of up to seven frames per second, for their shorter range settings, and the skew for speeds of up to 4 knots in many cases falls within the range resolution of the sonar. For mechanically scanned sonars the algorithm is limited to slow moving vehicles, typically much less than 1 knot, or vehicles equipped with a suitable dynamics model, or sufficient dead-reckoning sensors.

4.2 Data Processing and Segmentation

Each sector image is low-pass filtered to remove the backscatter noise. In this implementation we use a median filter, however other smoothing filters can also be implemented. The size of the window of the median filter will depend on the sector size, range, angular and range resolution and the gain setting of the sonar. Our procedure consists in manually calibrating the algorithm once for each parameter

setting. This procedure, although cumbersome, generally provides far better results.

Once the images have been smoothed the significant returns are obtained by applying a double threshold. The values of which are obtained from the histogram of the image. The double threshold uses eight-connectivity to define the neighbourhood of the pixels. The output from the double threshold is a binary image.

The algorithm subsequently extracts features for each observation. These features are the observations relative center of mass, with respect to the vehicle, the size in pixels of the target and the targets' first invariant moment [13]. The size of the targets are subsequently converted to meters squared. These features are the observations which will be fed onto the MHTF data association algorithm. Observations which are adjacent to the image edges are ignored as their center of mass will not be the true center of the object.

4.3 The MHTF Implementation

The proposed implementation uses both *clusters* and *super-clusters* to reduce the number of hypotheses [15]. A cluster is formed for a measurement that falls within the validation gate of an established target, any subsequent measurements that fall within that gate will belong to the same cluster. The validation gate is a threshold γ and the measurement must satisfy the following criteria,

$$v_i^T \mathbf{S}_i^{-1} v_i \leq \gamma \quad (14)$$

where the value for γ is obtained from the χ^2 distribution. If a measurement falls inside two clusters, those clusters will form a supercluster. Hypotheses will be formed such that measurements will only be associated to targets in belonging to the same cluster or supercluster. This system allows for a dramatic reduction in the amount of computation required in finding the most probable cumulative event.

4.4 The Stochastic Map Implementation

The stochastic map, as explained in section 3.1, holds the state of the vehicle and targets. In our application, we assume no knowledge of the vehicle's dynamic model, and no inputs from dead-reckoning sensors. Under these conditions the stochastic map will only be observable with at least two targets in view. Given these restrictions section 4.4.1 describes the assumed model. Section 4.4.2 describes the targets state vectors and the procedure for adding new targets onto the map.

4.4.1 Vehicle Model

Given no *a priori* knowledge of the vehicle model we assume a linear model description. This model has been found

to work well and can be used with any vehicle. The state of the vehicle takes the following form,

$$\mathbf{x}_v(k) = [x \quad \dot{x} \quad y \quad \dot{y} \quad \phi \quad \dot{\phi}]^T \quad (15)$$

with the following dynamic model,

$$\mathbf{F}_v(k) = \begin{bmatrix} \mathbf{F}_{v_x}(k) & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{F}_{v_y}(k) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{F}_{v_\phi}(k) \end{bmatrix} \quad (16)$$

where

$$\mathbf{F}_{v_x}(k) = \mathbf{F}_{v_y}(k) = \mathbf{F}_{v_\phi}(k) = \begin{bmatrix} 1 & dt \\ 0 & 1 \end{bmatrix} \quad (17)$$

And process noise,

$$\mathbf{Q}_v(k) = \begin{bmatrix} \mathbf{Q}_{v_x}(k) & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_{v_y}(k) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{Q}_{v_\phi}(k) \end{bmatrix} \quad (18)$$

where,

$$\mathbf{Q}_{v_x}(k) = \begin{bmatrix} \frac{1}{4} dt^4 & \frac{1}{2} dt^3 \\ \frac{1}{2} dt^3 & dt^2 \end{bmatrix} \sigma_{v_x}^2 \quad (19)$$

4.4.2 Target Model and Addition of New Targets

Each target has a state vector with states for its position, size and first invariant moment,

$$\mathbf{x}_i(k) = [x_i \quad y_i \quad s_i \quad m_i]^T \quad (20)$$

The updating of the state is done according to the procedure outlined in section 3.1.

Observations that were not associated to an existing feature will be added to the stochastic map state and covariance. The new observation $\mathbf{z}_{\text{new}} = [r \quad \theta \quad s \quad m]^T$ is estimated with respect to the vehicle's reference frame,

$$\mathbf{x}_{n+1}(k) = \begin{bmatrix} x_v(k) + r \cos(\phi + \theta) \\ y_v(k) + r \sin(\phi + \theta) \\ s \\ m \end{bmatrix} \quad (21)$$

The new map state and associated covariance will be

$$\mathbf{x}(k) \leftarrow \begin{bmatrix} \mathbf{x}(k) \\ \mathbf{x}_{n+1}(k) \end{bmatrix} \quad (22)$$

$$\begin{aligned} \mathbf{P}_{n+1 \ n+1}(k) &= \mathbf{L}_{\mathbf{x}_v} \mathbf{P}_{v \ v}(k) \mathbf{L}_{\mathbf{x}_v}^T + \mathbf{L}_{\mathbf{z}_{\text{new}}} \mathbf{R}(k) \mathbf{L}_{\mathbf{z}_{\text{new}}}^T \\ \mathbf{P}_{n+1 \ v}(k) &= \mathbf{P}_{v \ n+1}^T(k) = \mathbf{L}_{\mathbf{x}_v} \mathbf{P}_{v \ v}(k) \end{aligned} \quad (23)$$

where $\mathbf{L}_{\mathbf{x}_v}$ and $\mathbf{L}_{\mathbf{z}_{\text{new}}}$ are the Jacobian of equation 21 with respect to the robot vehicle state $\hat{\mathbf{x}}_v$ evaluated at $\hat{\mathbf{x}}_v(k)$ and to the new observation \mathbf{z}_{new} evaluated at \mathbf{z}_{new} .




Figure 2: Stochastic map and a sonar frame

5 Results

This section shows results for two experiments. The first experiment was carried out in the laboratory's tank. This experiment, section 5.1, was performed under controlled conditions and ground truth data is available for validation. The data for the second experiment, section 5.2, was recorded on a field trip, there is no ground truth data for the vehicle's or targets' position.

5.1 Tank Experiment

This experiment was carried out using a Tritech Dual Frequency Sonar. This is a mechanically scanned sonar and the frame rate is of two seconds. It has a horizontal beam-width of 2° and a vertical beam-width of 20° when operating at 675 kHz. The operating range was set at 5m, the sonar offers a 0.05 m range resolution. The sonar was mounted on the laboratory's planar Cartesian robot, this system allows for the sonar to be placed anywhere within the tank. The planar Cartesian robot has optical encoders allowing for a position accuracy of 1 mm. In this experiment we placed two cylinders in the tank and used them as targets. A heuristic was added to the algorithm so that it would ignore the tank walls and any measurements that fell outside. Figure 2 illustrates a frame of the sequence used in the experiment and the stochastic map, including the vehicle's trajectory at the end of the run and the tracked targets(numbered). The consistency of the map can be corroborated by examining figure 3, this figure shows the error in X and Y coordinates and the one standard deviation uncertainty bounds.

5.2 Field Experiment

The data for the experiments was obtained on trials at Oban on the west coast of Scotland. The sonar used was the SeaBat 6012 [17]. This sonar has a sector size of 90° by 15° . The sonar head contains all solid-state electronics required to form and transmit pulses at 455 kHz and receive returned energy into 60 1.5° electronically formed beams.




Figure 3: Innovation and uncertainty bounds

The sonar range was set at 10 meters. The sonar was carried by a diver towards a set of pier legs. No ground truth was available. The sequence consists of 95 frames, figure 4 shows the first and last frames respectively. The update rate of the sonar was of five frames per second.




Figure 4: Oban sequence: first and last frame

Given the lack of ground truth, the filter consistency is observed by plotting the innovation error and the one standard deviation uncertainty bounds. Figure 5 illustrates the results for target three. A qualitative consistency analysis is also possible by combining the last frame, subject to a rotation and translation, with the first frame. Figure 6 illustrates the outcome for such mosaic and the stochastic map ob-




Figure 5: Innovation and uncertainty bounds




Figure 6: Stochastic map and manual mosaic

tained with this data set. The map illustrates the trajectory followed by the sonar, the position of the tracked targets (numbered) and the uncertainty of both the position of the sonar and the targets, represented by ellipses. The mosaic was created manually and the aim was to minimise the error obtained from the greylevel values of corresponding pixels. The mosaic suggests that the heading difference between the first and last frames is of 32° clockwise and the translation is of 0.51 meters along the x-axis and 1.29 meters along the y-axis. These values can be compared with the output values of the algorithm which suggest an clockwise rotation of 30.14° and a translation of 0.34 meters along the x-axis and 1.29 meters along the y-axis. The error falls well within the standard deviation bounds of the filter which are 2.60° for heading and 0.19 meters and 0.10 meters for the translation along the x and y axis respectively.

6 Conclusions

The potential of the feature extraction and data association algorithm has been demonstrated with results using real data. The stochastic map offers accurate position fixes of both the obstacles and the vehicle itself. However the MHTF module is expensive in terms of computing power, exponentially increasing according to the number of targets being considered by the hypothesis matrix. Future work will compare the method with a nearest neighbour algorithm aided by simple heuristics and improved feature extraction techniques.

The results shown should be expected to improve as sensors and a valid vehicle model are integrated into the system. Further testing is planned using a mechanically scanned sonar on-board RAUVER, a house-built ROV. These future experiments will be carried out with the help of dead-reckoning sensors, so that the dependency on large feature densities is reduced.

7 Acknowledgments

This research is being supported by the European community under the MAST project CT97-0083 (ARAMIS).

References

- [1] D. B. Reid, "An algorithm for tracking multiple targets," *IEEE Transactions on Automatic Control*, 1979.
- [2] J. J. Leonard and H. F. Durrant-Whyte, *Directed sonar sensing for mobile robot navigation*, Kluwer Academic Publishers, 1992.
- [3] S. Thrun, D. Fox, and W. Burgard, "Probabilistic mapping of an environment by a mobile robot," in *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, Leuven, Belgium, May 1998, IEEE, vol. 2, pp. 1546–1551.
- [4] J. A. Castellanos, J. M. Martínez, J. Neira, and J. D. Tardós, "Simultaneous map building and localization for mobile robots: A multisensor fusion approach," in *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, Leuven, Belgium, May 1998, IEEE, vol. 2, pp. 1244–1249.
- [5] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," in *Autonomous Robot Vehicles*, I. Cox and G. Wilfong, Eds. Springer-Verlag, 1990.
- [6] J. K. Uhlmann, *Dynamic Map Building and Localization: New Theoretical Foundations*, Ph.D. thesis, University of Oxford, 1995.
- [7] H. J. S. Feder, *Simultaneous Stochastic Mapping and Localization*, Ph.D. thesis, Massachusetts Institute of Technology, 1999.
- [8] P. Newman, *On the Structure and Solution of the Simultaneous Localisation and Map Building Problem*, Ph.D. thesis, The University of Sydney, Sydney, Australia, 1999.
- [9] D. M. Lane, M. J. Chantler, and D. Dai, "Robust tracking of multiple objects in sector-scan sonar image sequences using optical flow motion estimation," *IEEE Journal of Oceanic Engineering*, vol. 23, no. 1, pp. 31–46, January 1998.
- [10] D. M. Lane, M. Chantler, D. Y. Dai, and I. Tena Ruiz, "Tracking and classification of multiple objects in multi-beam sector scan sonar image sequences," in *Proceedings of the 1998 International Symposium on Underwater Technology*, Tokyo, Japan, April 1998, pp. 269–273.
- [11] I. Tena Ruiz, Y. Petillot, D. M. Lane, and J. M. Bell, "Tracking objects in underwater multibeam sonar images," in *IEE Colloquium on Motion Analysis and Tracking*, London, UK, May 1999, pp. 11/1–11/7.
- [12] E. Trucco, Y. Petillot, I. Tena Ruiz, C. Plakas, and D. M. Lane, "Feature tracking in video and sonar subsea sequences with applications," *Computer Vision and Image Understanding*, no. 79, pp. 92–122, 2000.
- [13] I. Tena Ruiz, D. M. Lane, and M. Chantler, "A comparison of inter-frame feature measures for robust object classification in sector scan sonar image sequences," *IEEE Journal of Oceanic Engineering*, vol. 24, no. 4, pp. 458–469, October 1999.
- [14] P. S. Maybeck, *Stochastic models, estimation, and control. Volume 2*, vol. 141 of *Mathematics in Science and Engineering*, Academic Press, 1982.
- [15] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*, vol. 179 of *Mathematics in Science and Engineering*, Academic Press, 1988.
- [16] M. W. M. G. Dissanayake, P. Newman, H. F. Durrant-Whyte, S. Clark, and M. Csorba, "A solution to the simultaneous localisation and map building (slam) problem," Tech. Rep. ACFR-TR-01-99, Australian Center for Field Robotics, University of Sydney, Australia, January 1999.
- [17] Reson, Inc., *Seabat 6012's Operator Manual*.