The Fusion of Large Scale Classified Sidescan Sonar Image Mosaics

S. Reed, I. Tena Ruiz *Member, IEEE*, C. Capus *Member, IEEE*, and Y. Petillot

***Abstract*— This paper presents a unified framework for the creation of classified maps of the seafloor from sonar imagery. This is a longstanding problem in underwater imagery with direct applications to remote sensing. It presents difficult challenges in photometric correction, classification, navigation and registration and image fusion, drawing from a wide range of techniques. The paper uses recently developed methods to pre-process the images in order to compensate for the beam pattern of sonar devices and the motion of the acquisition platform. In a second step, the corrected images are classified using texture descriptors and standard classifiers. In parallel, the navigation of the sonar device is processed using Kalman filtering techniques. A Simultaneous Localisation And Mapping (SLAM) framework is adopted to improve the accuracy of the navigation. The classified images are fused within a Markovian framework. Two fusion models are evaluated. The first model uses a voting scheme regularized by an isotropic Markov Random Field when the reliability of each information source is unknown. The Markov model is also used to inpaint the regions where no final classification could be reached using pixel level fusion. The second model introduces formally the reliability of each information source into a probabilistic model. The performances of the two models are evaluated on synthetic images and real data. Finally, the system is demonstrated on a real mission containing hundreds of images and large scale classification maps are generated. It is shown that the Markovian fusion framework improves significantly the quality of the resulting classified maps.**

***Index Terms*— Registration, Classification, fusion, Markov Random Fields, side-scan sonar, mosaicing**

I. Introduction

Recent advances in the fields of underwater technologies and robotics have led to the development of Autonomous Underwater Vehicles (AUVs). The development of stable platforms [1]–[3], fitted with high resolution sonars, and capable of swimming close to the seabed has opened up the oceans to rapid and high resolution mapping, generating large volumes of data. These data have many practical scientific uses including reef management [4], oyster management [5], trawling impact assessment [6] and mine-counter measures [7]. This paper tackles the difficult problem of generating large scale accurate maps of the seabed from sonar imagery obtained using these platforms.

A large body of work on image based classification of sonar data exists, but the methods used have been based on single sonar images. The generation of large scale maps raises the problem of image registration, the difficult related problem of autonomous navigation in an environment deprived of GPS (Global Positioning Systems) and the generation of fused maps from multiple overlapping classified images which may contain contradictory information.

A. Sensors

In water, acoustics provides the main sensing modality. Electromagnetic waves attenuate rapidly and their operating range is limited to well below five meters in normal operating conditions. Sonar (SOund Navigation And Ranging) offers a good alternative providing the user with accurate resolution and long ranges of hundreds of meters. For imaging, side-scan sonar (SSS) and the emerging synthetic aperture sonar (SAS) provide very high resolution images of up to centimetric accuracy at up to 300 meters. These systems use the principle of a long antenna to generate a narrow acoustic beam [8], [9]. The beams illuminate a narrow stripe of the seabed at any one time. As the system moves through the water, towed by a survey vessel or mounted on an AUV, it generates a wide-area image of the sea bed (obtained as a concatenation of successive stripes) as shown in Fig. 2. The main parameters affecting the resolution of the generated images are the length of the antenna and frequency of the acoustic wave used (these determine the across track resolution) and the speed of the platform (AUV or towed body) which determines the along track resolution.

B. Image Formation Process and Pre-processing

The interpretation of sidescan imagery is a skilled procedure [8], [9]. There are many parameters of the image formation process contributing to intensity variations in recorded data which are quite separate from the influences of variations in seabed properties and textures. These are generally well understood and good models exist which can aid the process of automated seabed classification [10], [11]. In the current work the sonar data are preprocessed to correct for the influences of the sonar beam pattern and time-varying gain (TVG) [12], [13]. This enables the use of simpler and faster classification algorithms which is particularly beneficial for application to large area surveys.

C. Classification

Given the vast quantities of data produced, fast classification algorithms are required to produce seabed class maps representing textural variations and areas of scientific interest. Many supervised techniques have been developed to tackle this problem. Neural Networks and parametric statistical classifiers have dominated the scene [14]–[19] and have been coupled with feature extraction measures including one-dimensional cepstral and spectral features [20]–[22], fractal analysis [23], spatial point processes [24], grey level run-length measures [17], [25] and co-occurrence matrices [26],

[27]. Fuzzy logic analysis [28] and autoregressive models [29] have also been investigated for seabed classification.

These image-based techniques consider each image extracted from the sidescan data in isolation. In this paper a variant of the power spectrum feature set [21] is used with a simple parametric classifier for rapid supervised classification of the individual sidescan sonar images. However, unlike previous techniques, the classification result for each image is fused with the results derived from all of the images covering the same area of seabed.

D. Registration and Mosaicing

In order to produce large scale classified maps of the seabed, it is first necessary to register the individual sidescan sonar images. Given the position of the sensor in the world for each sidescan beam, it is possible to produce a geo-referenced image of the seabed. This process is called mosaicing [30], [31]. The quality of the mosaics produced will ultimately depend on the precision of the position information of each sonar beam, i.e. the vehicle's navigation precision. Underwater, navigation is a difficult problem as conventional GPS receivers do not operate. Underwater platforms therefore rely on dead-reckoning sensors which drift over time. To fix this drift, we propose to use a Simultaneous Localisation And Mapping (SLAM) technique based on the stochastic map developed for indoor robotics [32], [33] and adapted to sonar imagery [31], [34]. This is a solution, based on the use of landmarks detected in the sonar imagery to help the navigation, which does not require the addition of additional sensors and does not interfere with the data acquisition constraints (stable platform, fixed altitude).

E. Fusion

During a typical survey of the seabed, multiple views of the same area are normally collected from different view points. The fusion of these views enables the generation of improved large scale classified mosaics from the individual classified sonar images. To date, very little work has been done in image fusion in the underwater domain [35], [36] and it has been limited to simple multi-sensor fusion. To the author's knowledge, there have been no publications concerning the fusion of underwater imagery to produce large scale classified mosaics of the seafloor.

Fusion of multiple sources of information is a well established research field. When the information sources produce the same type of measurements, standard fusion techniques such as Bayesian theory [37], Fuzzy Logic [38] and Dempster-Shafer theory [39]–[41] can be used. However, for classifiers which consider different types of input measurements or features, it is often not possible to consider the computed output measurements to be estimates of the same posterior probability [42], making fusion difficult. When little is known about the information sources, or when they produce information at a high level of abstraction, voting schemes can be successfully used [25], [43]. This is particularly appealing for underwater imaging systems which currently favor 'black box' approaches to classification.

Fusion for classification within the image domain allows contextual information to be considered and to date has been mainly applied to remote sensing [37], [38], [44]. The fusion of multiple images is generally performed at the pixel level [41], [45]–[48], allowing information from surrounding pixels to be considered when classifying each pixel. An effective method for incorporating this spatial information is the use of Markov Random Fields(MRF) [49]–[51].

This paper details a fusion model for fusing registered, classified sidescan sonar images of the seafloor. Fusion is conducted at the pixel level where each classifier outputs a single class decision for each pixel. This maintains the generality of the fusion scheme and allows classified images from different classifiers to be fused together.

Two separate models are presented: the first model uses a voting scheme to initialize the fused class map. After, a MRF model is used to incorporate contextual information, smoothing the final result and 'inpainting' regions of pixels which are *unclassified* after the Voting process. This ensures all pixels within the image, for which there is data, are successfully classified within the final fused result. The standard Markovian prior term is altered so that pixels labeled as *unclassified* from the Voting scheme, do not contribute to the Markovian probability. This ensures that only pixels which have been successfully fused and classified contribute to the fusion process for the other pixels. The second model considers the case when the reliability of the classifier is known [52]. This information is stored using confusion matrices and allows the Voting Scheme from the first model to be replaced by a probabilistic, Markovian framework.

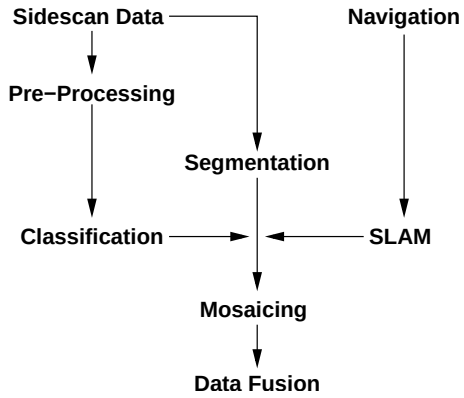

Fig. 1. **Data flow for formation of fully classified fused sidescan mosaics**.

Fig. 1 illustrates the progression from the raw sensor and navigation data through to the fully classified fused sidescan image mosaics resulting from the processing steps outlined in this paper.

F. Layout

The image formation mosaicing and pre-processing of the sidescan sonar data are discussed in Section II. Supervised classification using features derived from the sonar swath power spectra is described in section III. Section IV describes the SLAM-RTS technique for improving and smoothing the

navigation solution. Section V details the two models presented for Fusion of the classified sidescan sonar images. Results are presented on simulated and MeasTex sample data to illustrate the differences between the two models and quantitatively measure the effectiveness of these approaches. Section VI contains the results of the Fusion model on real, classified sidescan sonar mosaics. Section VII concludes the paper and outlines future research plans.

II. Side-scan Sonar Image Formation and Data Pre-processing

A. Image Formation

A sidescan sonar is an acoustic device comprising a long horizontal antenna. This antenna forms a narrow beam orthogonal to the antenna, along the vehicle track, and a wide beam across the track. As the antenna is moved forward, ideally in a straight line, the beam generated covers the seabed uniformly. The geometry of the sidescan image formation process is represented in Fig. 2. A sonar image is generated by concatenating these beams into a 2D image as seen in Fig. 3 (a). It is important to understand that such images are not in Cartesian coordinates but in time coordinates. The x coordinates correspond to the time at which the beam was emitted from the sonar while the y coordinates corresponds to the time of flight of the pulse in the across track direction. Converting these time-time images to images in real-world coordinates is called geo-referencing or mosaicing and is critical to our applications. A common feature in all raw sidescan imagery is the largely black region in the centre of the image. This corresponds to the transit time of the acoustic wave through the water before reflection from the seabed and indicates the altitude of the sensor. More details on the image formation process can be found in [9].

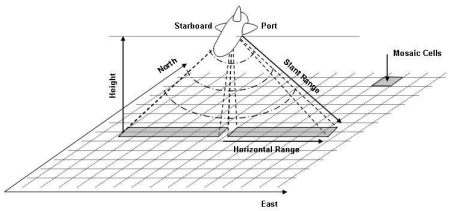

Fig. 2. **Diagram Showing the Assumed Side-scan Sonar Geometry**.

B. Preprocessing

In many emerging applications low altitude surveys are required. In such cases even quite small changes in vehicle altitude can affect the sonar image dramatically. Prior to classification the image data used here have been preprocessed using an advanced radiosity correction algorithm [12]. This is useful because it treats purely range-dependent artifacts, such as residual TVG effects separately to angular effects such as the influence of the sonar beam pattern. Separate correction factors are calculated for each. Whilst this gives better performance than standard radiosity correction algorithms in the presence of sensor altitude changes, platform stability is still assumed with respect to pitch and roll.

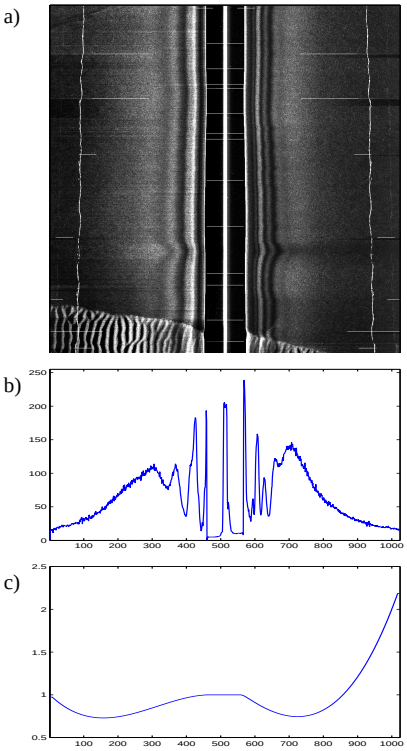

Fig. 3. **(a) sample raw image (1000 scan lines); (b) beam pattern estimate; (c) residual TVG estimate. Differences in ranges for y-axes derive from the methods of calculation. Values in (b) are referenced to a target grey level, hence range 0–255, values in (c) are referenced to a level of 1.0.**

A sample raw image and the estimated beam pattern and residual TVG profile for these data are shown in Fig. 3. The complexity of the beam pattern is apparent with four significant lobes in the port channel and as many as six in the starboard channel. The differences in the scales of the y-axes for the beam profiles and TVG estimates result from their methods of calculation and application within the radiosity correction algorithm [12].

The corrected image is shown in Fig. 4. In some places the beam pattern correction has failed, as indicated by the white arrow. This arises from the behavior of the vehicle, which rolls on turns. With each course adjustment the small degree of roll affects the symmetry of the beam pattern on the seabed, so that it is poorly compensated near the water column. At these points classification accuracy is affected. However, where there are overlapping images, the data fusion techniques described below compensate well for any resulting misclassification.

C. Mosaicing

The sonar mosaic algorithm used in this paper assumes the geometry shown in Fig. 2. Under this assumption each sonar channel (port and starboard) insonifies a rectangular area on the sea floor. The length of the rectangle is determined by the slant range of the sonar (the maximum range of the sonar) and the height of the vehicle, assuming a flat seabed. The equation for this length, referred to as horizontal range r_h is:

$$r_h = \sqrt{r_s^2 - h^2} \tag{1}$$

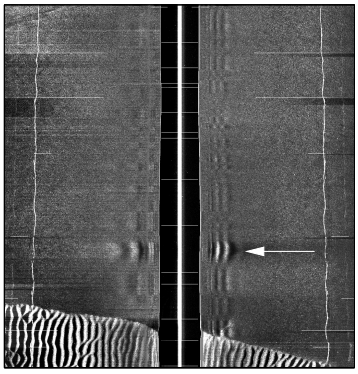

Fig. 4. **Corrected image. The correction algorithm cannot compensate for changes in sensor attitude, such as roll on trajectory corrections, see arrow above.**

where the slant range r_s is obtained from:

$$r_s = T \times c/2 \tag{2}$$

with c the speed of sound in the water, assumed constant, and T the time of the last return on each beam. The height of the vehicle h is found using:

$$h = t_b \times c/2 \tag{3}$$

where t_b is the time to the first return of the sonar beam. The width of the beam is determined by the spacing between two successive beams, this guarantees that all the mosaic cells between beams are filled.

The seabed is represented as a flat two dimensional grid of mosaic cells (see Fig. 2). The resolution of the mosaic cells can be set by the operator. Using simple geometry both channels can be superimposed if the orientation and position of the sonar is known. Each mosaic cell will take the value of the intensity cell on the beam closest to it. In this paper if a mosaic cell takes more than one value then, in the case of the pre-processed image mosaics, the data will be averaged or, in the case of the classified data, the cell will be left as unclassified.

An example output from the mosaicing algorithm is given in Fig. 5 which shows a mosaic obtained from geo-referencing the data from Fig. 4.

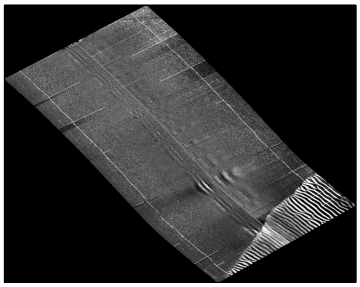

Fig. 5. **Sample side-scan mosaic.**

III. Classification of Side-scan Data

Three seafloor textures have been identified for segmentation of the sidescan data, defining three classes: flat sediments, sand ripples, and complex regions. Suppression of the beam pattern effects and some of the residual TVG effects, as described in §II-B above, improves the images to the point where a fast supervised classification scheme can be combined with a relatively simple, easily generated feature set.

The features used are based on Pace and Gao's frequency based sediment classification scheme [21]. In the current data the classes have relatively large scale textures, which confines much of the discriminatory content to relatively low frequency bands. Overlapping 64-sample Gaussian windowed FFTs are used to generate the one-dimensional power spectra and this allows for identification of changes in texture across the sonar swath.

If $s_j(t)$ represents a single line of sonar data and $W_i(t)$ represents the Gaussian window centered at position i, the normalized power spectrum can be defined by,

$$\begin{aligned} P_{i,j}(f) &= |F[W_i(t)s_j(t)]|^2 \\ \tilde{P}_i(f) &= \frac{1}{n}\sum_{j=1}^{n} P_{i,j}(f) \\ P_{iN}(f) &= \tilde{P}_i(f) \Big/ \int_0^{f_{\max}} \tilde{P}_i(f)df \end{aligned} \tag{4}$$

Spatial frequency bands within the normalized power spectra are identified which give a good separation between the classes. The training set used for these data comprised three small sections, one for each texture, of 200×200 pixels. The training images extracted from the full dataset are shown in Fig. 6.

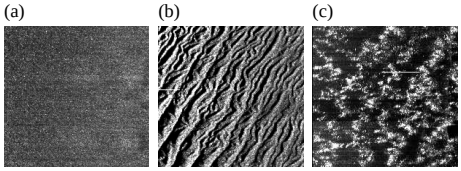

Fig. 6. **Training data. Three small images extracted from the full dataset. Classes are: (a) flat sediment; (b) sand ripples; (c) complex texture.**

The averaged normalized spectra for the three training samples are presented in Fig. 7. The bimodal nature of the spectrum for the sand ripples class occurs because there are two dominant scales for the sand ripples. The training image used has been chosen to be representative of both the large and small scale ripple textures. It is also possible with this data set and classification scheme to separate these two ripple textures and define a four-class training set.

Returning to the three-class problem, three features are defined by the crossing points of the averaged normalized spectra derived from the training data. These give the proportion of the spectrum lying in sample bands 1-4, 4-12 and 16-32, corresponding to crossing points lying at $f_{\max}/8$, $3f_{\max}/8$ and $f_{\max}/2$.

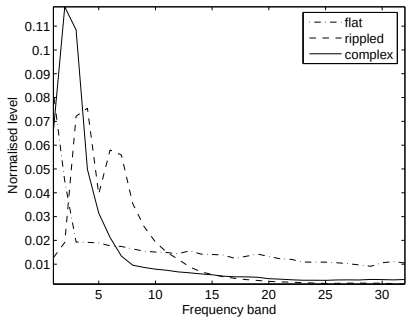

Fig. 7. **Averaged normalized power spectral densities for the three training sets. Three features were defined with band limits at 1-4; 4-12; 16-32. The rippled texture is characterized by two dominant peaks corresponding to large and small scale ripples which dominate this region.**

$$
\begin{aligned}
D_{f1} &= \int_{1}^{f_{\max}/8} P_{iN}(f) \Big/ \int_{1}^{f_{\max}} P_{iN}(f) \\
D_{f2} &= \int_{f_{\max}/8}^{3f_{\max}/8} P_{iN}(f) \Big/ \int_{1}^{f_{\max}} P_{iN}(f) \\
D_{f3} &= \int_{f_{\max}/2}^{f_{\max}} P_{iN}(f) \Big/ \int_{1}^{f_{\max}} P_{iN}(f) \qquad (5)
\end{aligned}
$$

In classifying a complete sonar image, the three features are generated from the averaged normalized spectral density formed from four successive lines of data. The same 64-sample sliding Gaussian windowed FFT is used and boundary problems between sonar channels are minimized by closing up the water column. This is done simply by shifting the scanlines on the assumption that there will generally be continuity in seabed textures between sonar channels.

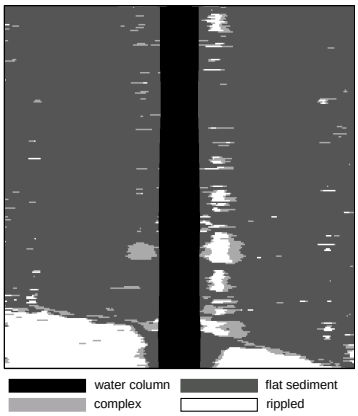

Fig. 8. **Classmap generated from the image introduced in Fig. 3. Water column – black; flat sediment – dark grey; complex – light grey; rippled – white. Errors are noted where the correction algorithm has failed due to vehicle attitude changes during course corrections and in the transition regions between textures.**

Fig. 8 shows the initial classification result for the image introduced in Fig. 3 above. Misclassifications are greatest near the water column where the correction algorithm has failed. There are some boundary errors, with pixels classified as complex texture in the transition region between flat sediment and sand ripples. Misclassification further from the water column is due primarily to incomplete elimination of the influence of the surface return and crosstalk from other sensors. These effects are particularly prominent in this data set and add to the difficulty of the classification task. These misclassifications can frequently be rectified by the proposed fusion scheme.

IV. Simultaneous Localisation and Mapping (SLAM)

In order to create an accurate mosaic of the classified maps, good navigation is crucial. Conventional GPS and Differential GPS (DGPS) receivers do not work underwater. Therefore, when submerged the side-scan sensor must be localized using dead-reckoning with depth sensors, Doppler Velocity Logs (DVL), Inertial Navigation Systems (INS) and/or compasses [53]–[55]. To correct drift on the dead-reckoning, the sonar must either be equipped with a GPS/DGPS system and surface intermittently to get a new fix, thus perturbating the data acquisition process, or it must be equipped with acoustic receivers capable of triangulating the position with respect to either acoustic beacons on the seabed, known as Long Base Line (LBL) navigation, or to acoustic beacons on a support vessel, known as Short Base Line (SBL) or Ultra Short Base Line (USBL) [56]. The costs associated with high performance INS systems and with setting up acoustic nets or mobilizing a vessel are considerable and new techniques have been sought to localize positions underwater. Terrain matching methods will use known maps of the environment and data from payload sensors to find the dead-reckoning drift [57], [58]. The purpose of SLAM is to build a map of the environment and use that same map to localize [32], [33]. Recently SLAM techniques have been developed to work with a side-scan sonar [31], [34]. This paper uses this method in order to geo-reference classified side-scan images. Previous work has demonstrated the potential of this method when fusing non-classified data using Gabor wavelets [59].

The data from the navigation sensors is fused in order to localize the side-scan sonar. The technique used to fuse the navigation data in this paper is the stochastic map smoothed using a Rauch-Tung-Striebel (RTS) fixed-interval smoother, it will be referred to as SLAM-RTS. The stochastic map keeps the estimates of the position and creates a map of landmarks to represent the environment. These landmarks are then used to aid localization of the vehicle. It is a SLAM method that works iteratively to provide an estimate of the position at the latest iteration. In order to improve the accuracy of the solution and to smooth it, post-processing is required. The next two sections provide a detailed look at the algorithms.

A. The Stochastic Map

The stochastic map is an augmented state Extended Kalman Filter (EKF) [60], [61]. It adds new states to the state vector to accommodate new landmarks as they are observed [62]. A typical stochastic map state vector is of the form:

$$\mathbf{x} = [\mathbf{x_v} \mathbf{x_1} \dots \mathbf{x_n}]' \qquad (6)$$

where $'$ is the transpose of a vector or matrix, $\mathbf{x_v}$ holds the state of the side-scan sonar and $\mathbf{x_1}, \ldots, \mathbf{x_n}$ hold the state of the n landmarks in the map.

The stochastic map also stores and maintains all the covariances and correlations between the states. Furthermore, it has been proved in [63] that, in the limit as the number of observations increases, the covariance associated with any single target location estimate is determined only by the initial covariance in the vehicle location estimate and, in the limit, all the target estimates become fully correlated. These properties make the stochastic map highly desirable. With fully correlated landmarks, an observation of any of the landmarks will help correct the whole map. It can also take advantage of the wealth of literature published on Kalman filters. The update equations of the stochastic map are the familiar EKF update equations. To propagate the state:

$$\hat{\mathbf{x}}_\mathbf{v}(k) = \mathbf{f_v}[\hat{\mathbf{x}}_\mathbf{v}(k-1), \mathbf{u}(k), \mathbf{0}, k] \tag{7}$$

where $\hat{\mathbf{x}}_\mathbf{v}(\cdot)$ is the side-scan sonar's estimated state and $\mathbf{f_v}[\hat{\mathbf{x}}_\mathbf{v}(\cdot), \mathbf{u}(k), \mathbf{0}, k]$ is its dynamic model. And its associated covariance will be propagated thus:

$$\mathbf{P}(k) = \mathbf{F_{x_v}}\mathbf{P}(k-1)\mathbf{F}^\mathbf{T}_\mathbf{x_v} + \mathbf{F_{w_v}}\mathbf{Q}(k)\mathbf{F}^\mathbf{T}_\mathbf{w_v} \tag{8}$$

where $\mathbf{F_{xv}}$ is the Jacobian of the dynamic model with respect to the side-scan sonar state, used to linearise the state of the side-scan sonar error $\tilde{\mathbf{x}}_\mathbf{v}(k-1)$, and $\mathbf{F_{w_v}}$ is the Jacobian of the dynamic model with respect to the process noise. The corrected state estimate becomes:

$$\hat{\mathbf{x}}(k+1) = \hat{\mathbf{x}}(k) + \mathbf{K_i}(k)\upsilon_\mathbf{i}(k) \tag{9}$$

where $\mathbf{K_i}$ is the gain of the filter and $\upsilon_\mathbf{i}$ is the innovation. Its associated covariance is updated according to:

$$\mathbf{P}(k+1) = \mathbf{P}(k) - \mathbf{K_i}(k)\mathbf{S_i}(k)\mathbf{K}^\mathbf{T}_\mathbf{i}(k) \tag{10}$$

where $\mathbf{S_i}(k)$ is the innovation covariance.

For more details on this implementation of the stochastic map, the interested reader should refer to [34].

B. SLAM-RTS

The Kalman filter and EKF use all measurements up to the last iteration to estimate the state at that last iteration. The RTS smoother uses all measurements *before* and *after* each iteration to estimate the state at each iteration [64]. It is a post-processing filter that works on the stored outputs of a Kalman filter by re-processing it. The smoother works by combining a forward pass Kalman filter with a backward pass filter. It was originally designed to work with fixed size state vectors. However, the stochastic map adds new states to the state vector as it observes new landmarks. The SLAM-RTS algorithm adapts the RTS fixed-interval smoother to work with the stochastic map by fixing the size of the state vector to the size of the stochastic map on the last iteration. The SLAM-RTS algorithm ensures numerical stability in matrix operations by adjusting the estimates of the landmarks' states and covariances at all iterations before they have been observed to the values when they are first observed. The output of the SLAM-RTS has been shown to improve the accuracy of the stochastic map solution [34], as well as providing trajectories more suitable for creating and superimposing mosaics [31], [59].

V. Fusion of Mosaiced Classification Data

This section presents two pixel level models for the fusion of multiple classified sidescan sonar mosaics. The first model fuses classified maps using a simple voting scheme while the second model integrates class reliability indexes in the fusion scheme. Both models are formulated within a multi-source Markovian framework to take advantage of contextual information and improve classification accuracy. The Markovian model is presented first as it is common to both outlined models. The details specific to each of the two models are presented later.

A. Markov Model for Image Fusion

Let us assume first that each of the input class maps is defined on a lattice S where label s specifies a specific pixel location. Two random fields $\mathbf{X}$ and $\mathbf{W}$ are defined. $\mathbf{X} = \{X_s, s \in S\}$ describes the classification field provided by each input map and $\mathbf{W} = \{W_s, s \in S\}$ describes the final fused classification map. For K input class maps, $X_s = (X_s^1, \ldots, X_s^K)$ takes its values from the finite set of classes $\Omega = \{\omega_1, \ldots, \omega_M, \alpha, \gamma\}$. The set Ω contains M recognized seafloor classes, the *unclassified* label α and the *unmeasured* label γ. Label α is allocated to $X_s^j, j \in \{1, \ldots, K\}$ when data is received regarding pixel s in image j but a classification based on the data provided is not possible. Label γ, *unmeasured*, is used when no data is received regarding pixel s, ensuring it is not possible to provide a classification X_s^j.

The distinction between classes α and γ is important during the fusion process. If all the input class maps provide input $X_s^j = \gamma, \forall j$, then there has been no data retrieved by the sensor regarding pixel s. Therefore, pixel s should not be classified by the fusion process. If all the input class maps provide input $X_s^j = \alpha, \forall j$ then while none of the input images have provided a classification for pixel s, the sensor has received data regarding this region of seafloor. In this instance, the Fusion model will return an output classification for pixel s by considering the pixel classifications within the neighboring region.

The fusion problem consists of estimating the true classified map $W = w$ from the individual classified maps $X = x$ where $x = x_s, s \in S$ are K classified maps of the same scene. The field $W = \{W_s, s \in S\}$ is said to be Markovian with respect to neighborhood $\eta = \{\eta_s, s \in S\}$ if its distribution can be written as

$$\begin{aligned} P_{W_s}(W_s = w_s | W_r = w_r, r \neq s) = \\ P(W_s = w_s | W_r = w_r, r \in \eta_s) \end{aligned} \tag{11}$$

This formalizes that $P(W)$ is a local probability and that the fused class label w_s for pixel s is dependant only on the class labels of the pixels within its neighborhood η_s. For simplicity, the fusion model described in this paper assumes a second order isotropic neighborhood. This neighborhood η_s for pixel

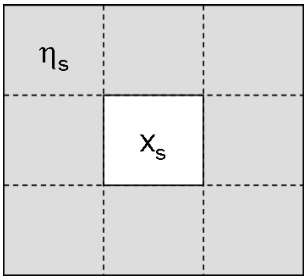

Fig. 9. **The Markovian prior assumed a second order isotropic neighborhood. The class of pixel s is therefore only dependant on its 8 nearest neighboring pixels.**

s can be seen in Fig. 9. Further reading regarding MRF models can be found in [50], [51], [65], [66].

The Markovian property field W allows Prior Probability $P(W)$ to be written in the form

$$P_W(w) = \frac{1}{Z}\exp(-E(w)) \tag{12}$$

Where Z is a normalizing constant and $E(w)$ is an energy term.

For the Fusion model, the problem of maximizing probability $P_W(w)$ can be re-cast to the local problem of maximizing energy

$$U(w_s) = \sum_{t\in\eta_s} \beta\delta(w_s, w_t)[1-\delta(w_t,\alpha)][1-\delta(w_t,\gamma)] \tag{13}$$

for pixel s. In equation 13, $\delta(.)$ is the Kronecker Delta symbol and β controls the importance of the Markovian prior. For all cases in this paper, $\beta = 1.0$. As can be seen from equation 13, neighboring pixels labeled as *unclassified* or *unmeasured* (and therefore do not contain any useful information to aid the fusion process) do not contribute to the Markovian prior. The minimization of $P(W)$ is performed using the Iterated Conditional Modes method [51]. In this method, a raster scan is used to iteratively visit all the pixels in field W. If $w_s = \gamma$ *unmeasured*, the pixel is not considered further and the pixel remains *unmeasured*. Otherwise, w_s is allocated to the class which locally maximizes $U(w_s)$. This method of segmentation produces a local maximum of $P(W)$. The ICM procedure is iterated until there are no pixel changes within a full image scan. The details of how the energy term described in equation 13 is applied within each of the two Fusion models is provided in the following sections.

B. The Voting/MRF Fusion Model

The Voting/MRF Fusion Model assumes that each image j provides a classification result for each pixel label x_s^j. The fusion field W is initialized by using an adaptation of the Generalized Majority Voting [43]. In this model, a summed binary function $T_s(\omega_i)$ for pixel s, and each recognized seafloor classes $\omega_i, 1 \le i \le M$ is specified as

$$T_s(\omega_i) = \sum_{j=1}^{K} \delta(x_s^j, \omega_i) \;\; for \; 1 \le i \le M \tag{14}$$

where the sum is over all the inputted class images and as before, $\delta(.)$ is the Kronecker Delta Function. This function is not specified for the *unclassified* or *unmeasured* classes.

The initial fusion Field W is then specified as:

$$\begin{aligned} w_s &= \gamma \; if \; \sum_{j=1}^{K} \delta(x_s^j, \gamma) = K \\ &= \tau \; if \; T_s(\tau) = \max T_s(\omega_i) \ge \frac{2}{3}K^s \\ &= \alpha \text{ otherwise} \end{aligned}$$

where K^s is the number of images which do not provide an *unclassified* or *unmeasured* classification for pixel s.

In this model, the adapted voting function first checks to see whether all the images provide an *unmeasured* classification. If they do, the pixel is classified as *unmeasured* within the Fusion Map. For w_s to be allocated a class which is not *unmeasured* or *unclassified*, class $\tau \in \{\omega_1, \ldots, \omega_M\}$ must have the largest summed binary function with a value greater than or equal to $\frac{2}{3}K^s$. If this rule is not met, the pixel is labeled as *unclassified*.

Once the voting rule specified in equation 15 has been used to initialize Fusion Field W, the Markov Energy term $U(w_s)$ in equation 13 is considered. The ICM technique described in section V-A is used to locally maximize $U(w_s)$ and complete the Fusion process.

Using the above fusion approach ensures several points:

- Pixels allocated as *unmeasured* in all the images are left so in the fused image as no measurements have been obtained regarding these pixels. They are not allocated a valid seafloor class.
- Pixels initialized as *unclassified* in Fusion Field W as a result of the Voting Scheme are always allocated a seafloor class during the ICM process. In effect, the *unclassified* regions are eroded away by the Markovian prior (inpainting). The Markov model allows these regions to be classified by considering the surrounding classified regions.
- The ICM process produces a smoothed version of the initialization produced by the Voting process.

C. The Probabilistic Fusion Model

The Voting scheme used to initialize the Fusion Class Map in the Voting/MRF model described in section V-B assumes that each input source is equally reliable. A more balanced fusion decision should also consider the reliability of each source. One possible representation of source reliability is to use confusion matrices.

Assume each source $j \in \{1, \ldots, K\}$ produces an input class map X^j as well as a Class Confusion matrix C^j. This matrix is obtainable through training (using a supervised system) and provides a measure of the likelihood term $P(x_s|w_s)$ where it is assumed that random variable X_s is independently conditional on W [65]. This assumption is used extensively in image-based Markov solutions. Using training data, these matrices can be estimated, specifying the probability of source j providing a $x = \tau$ classification decision given that it is

known that the actual classification of the decision is $w_s = \epsilon$. For pixel s and input source j,

$$P(x_s^j = \tau | w_s^j = \epsilon) = C_{\epsilon\tau}^j \quad (15)$$

We assume that the Confusion Matrix C considers only recognized seafloor classes $\omega_1, \ldots, \omega_M$. Assuming the input sources are independent (we assume the random variables $X_s^1, \ldots, X_s^K$ are independent conditionally on W [47]), we can write that

$$P(x_s^1 = \tau, \ldots, x^K = \sigma | w_s = \epsilon) = C_{\epsilon\tau}^1 \times \ldots \times C_{\epsilon\tau}^K \quad (16)$$

Using the likelihood term expressed in equation 16, it is possible to initialize the Fusion Field W. Instead of initializing the field using a voting scheme where each source is considered equally reliable, the model now considers each inputs reliability. As before, if all the sources provide a $x_s = \gamma$ or α decision, w_s is initialized as *unmeasured* or *unclassified* respectively. Otherwise the label field W is initialized by considering equation 16. Ten samples are drawn from the class $w_s = \epsilon$ probability distribution which maximizes the likelihood value described in equation 16. Pixel s is initialized with class w_s if the sampled probability compares favorably to a random number in the range [0,1] for a majority of the samples. If this criterion is not met, the pixel is labeled as *unclassified*. This initialization technique is analogous to the generalized majority technique used in the Voting/MRF Fusion model.

After initialization, an iterative process is again carried out to complete the fusion. Pixels are visited randomly and allocated a classification which minimizes posterior energy

$$\begin{aligned} U^{post}(w_s, x_s^1, \ldots, x_s^K) = \\ z_s[1 - \delta(w_s, \alpha)] + \\ \beta \sum_{t \in \eta_s} \delta(w_s, w_t)[1 - \delta(w_t, \gamma)][1 - \delta(w_t, \alpha)] \end{aligned}$$

where as in the Voting/MRF model, pixels initialized as $w_s = \gamma$ are not considered or changed. As before, $\delta(.)$ is the Kronecker Delta function and

$$z_s = \ln[P(x_s^1, \ldots, x_s^K | w_s)] \quad (17)$$

As can be seen from equation 17, the likelihood term is considered only if the pixel is currently not classified as *unclassified*. The Markovian term is the same as the one discussed in equation 13 for the Voting/MRF model. Pixels are visited randomly rather than using the more deterministic raster scan to allow more mixing between the classes. The number of pixel visitations will affect the quality of the final result. The results provided here used $4N$ pixel visitations where N is total number of pixels in field W.

The probabilistic fusion process is completed by applying Energy term $U(w_s)$ in equation 13 to field W in an iterative manner, again using a raster scan. This was done to again ensure the final Fusion Classification contained no *unclassified* pixels. Unlike the raster scan approach, randomly visiting the pixels does not ensure that all initially *unclassified* pixels are classified in the final result. The final raster scan considering only $U(w_s)$ ensures that any remaining *unclassified* pixels are given a recognized seafloor classification in the final, fused map.

Classifier	% Classification Accuracy
C1	88.65
C2	75.90
C3	74.91
C4	48.66
C5	49.94

TABLE I

THE CLASSIFICATION ACCURACY OF THE 5 CLASSIFIERS USED TO CLASSIFY THE MEASTEX TEXTURE IMAGE SHOWN IN FIGURE 10

D. Results

1) Comparison of the Models on MeasTex Images: This section evaluates the Voting/MRF and Probabilistic fusion model on an image comprised of MeasTex [67] texture samples. Classification results have been obtained using 5 different classifiers. The first is a parametric linear discriminant classifier using co-occurrence matrix features(C1). The second and third consider fractal features using a non-parametric kNN classifier(C2) and a parametric linear discriminant classifier(C3) respectively. The fourth and fifth classifiers consider frequency based features, again with the kNN classifier(C4) and the linear discriminant classifier(C5) respectively. All 5 classifiers are supervised systems. Therefore the confusion matrices required for the Probabilistic model can be obtained from the training data. The test image and the corresponding ground truth result can be seen in Figure 10.

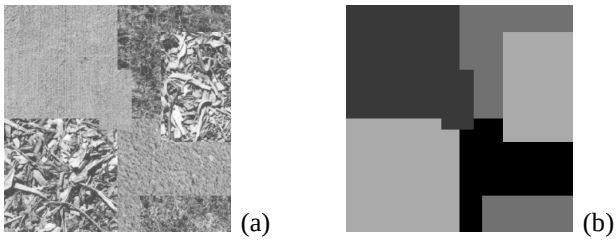

(a) (b)

Fig. 10. **(a) A texture image composed of 4 MeasTex textures. (b) The ground truth of the textured image displayed in (a).**

The performance accuracy for each of the 5 classifiers operating in isolation can be seen below in Table V-D.1.

The classification results for 3 of the classifiers (C1, C2 and C4) are shown in Figure 11 (C3 and C5 are visually quite similar to C2 and C4 respectively). The figure also contains the classification result obtained from fusing the C1, C2 and C4 classifier results using the Probabilistic model. This has been included for comparison purposes.

The classification accuracy obtained from fusing the different classifiers using both the Voting/MRF and the Probabilistic model can be seen below in Table V-D.1.

Table V-D.1 shows the Probabilistic model outperforming the MRF/Voting model in all cases as expected. The fusion

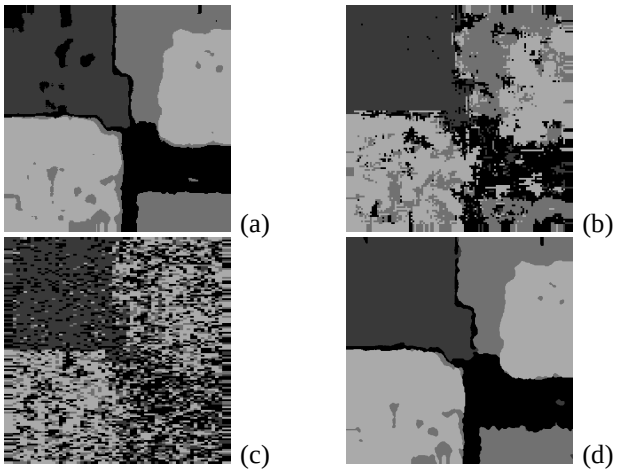

Fig. 11. **(a) Classification result from (a) classifier C1 (b) Classifier C2 (c) Classifier C3 (d) Fusing C1, C2 and C4 using the Probabilistic model.**

Classifiers Fused	Voting/MRF Accuracy	Probabilistic Accuracy
C1 C2 C3	85.58	91.72
C1 C2	91.29	92.56
C1 C3	89.51	92.09
C2 C3	80.86	82.47
C1 C2 C4	89.08	92.72
C4 C5	62.94	63.50

TABLE II

THE CLASSIFICATION ACCURACY OBTAINED FROM FUSING VARIOUS CLASSIFIERS USING BOTH THE VOTING/MRF AND PROBABILISTIC MODELS.

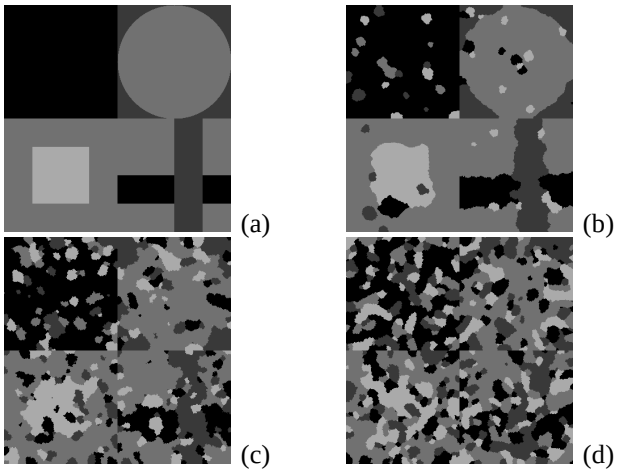

Fig. 12. **(a) The ground truth of the simulated image used to test the fusion models. (b)-(d) Simulated classification results from a 90%, 70% and 50% accurate classifier respectively.**

Input Image Accuracy %	Voting/MRF Accuracy %	Probabilistic Accuracy %
100	99.73	99.75
90	96.70	96.97
80	93.51	94.20
70	91.92	91.94
60	84.21	85.32
50	75.11	76.14

TABLE III

THE CLASSIFICATION ACCURACY OBTAINED FROM FUSING 4 SIMULATED CLASSIFICATION RESULTS OF A GIVEN CLASSIFICATION ACCURACY USING BOTH THE VOTING/MRF AND PROBABILISTIC MODELS.

operations containing C2 and C3 can be seen producing lower classification results than is perhaps expected from considering 2 good classifiers (the Voting/MRF fusion of C1, C2 and C3 is the only example where the final fused result is lower than any of the classifiers considered in isolation). This is because C2 and C3 use the same features and so generally classify and miss-classify the same image regions. Any miss-classifications present in either of these results are simply reinforced by the second. This lends strength to the argument that fusion is most effective when considering results obtained from different sources. This can be seen in the high classification accuracy obtained from fusing C1, C2 and C4, even though C4, when considered in isolation, is a poor classifier.

2) Comparison of the Models on Synthetic Images: This section evaluates the Voting/MRF and Probabilistic fusion models on a simulated example. A simulator model was used to produce classification results with a given classification accuracy, using the same ground truth image. Figure 12 contains the ground truth image used in this section as well as a simulated result from hypothetical 90%, 70% and 50% efficient classifiers. As can be seen from Figure 12, the simulator was developed to produce *regions* of miss-classified pixels, as would be observed from a real classifier.

Simulated classifier output for the ground truth image shown in Figure 12 was produced for a range of classifier accuracies. Four different classified images for each tested input classification accuracy were produced. The four images were fused using the Voting/MRF and Probabilistic models. The simulator was set up to ensure that any miss-classified pixels were spread evenly over the other classes. The results obtained by the 2 tested models to fuse the 4 images for each classification accuracy value, can be seen in Table V-D.2.

As Table V-D.2 shows, both fusion models generally produce a classification result which is better than any of the input images considered in isolation (the exception is the perfect input classifiers where the MRF component of the models has smoothed the final result). The Probabilistic model can be seen to outperform the Voting/MRF model in all cases, although both provide good results in all of the examples, even when the input classification accuracy drops to 50%. Further results could be provided using varying numbers of input images and different classification accuracies. However, this study lies out with the scope of this paper. Section V-D has demonstrated that both of the fusion models presented produce a final classification result that is generally higher in accuracy than any of the input images considered individually.

VI. CREATING LARGE SCALE MOSAICS

The following results were obtained by processing data gathered during the BP'02 experiments carried out by the SACLANT Undersea Research Centre in La Spezia, Italy. The side-scan data was gathered by a REMUS AUV [3].

The navigation output from REMUS and the side-scan data have already been used to create large scale mosaics of the observed region [68]. The same navigation solution has been used to geo-reference the data. The AUV mission lasted for 2 hours, 57 minutes and 8 seconds. It followed a set of parallel, regularly spaced and overlapping linear tracks (typical for rapid environmental assessment missions).

The data have been classified using the techniques outlined in section III.

A. Fusion

The fusion of the classified mosaics was carried out using the Voting/MRF model. This is because no ground truth information was available on the data set with which to obtain classifier reliability information.

Fig. 13 shows four overlapping mosaics of sector 1 (a pre-defined region) created by geo-referencing 17 linear tracks. All of the resulting mosaics created using the overlapping tracks are geo-referenced to the same reference frame (sector 1) and will constitute an input to the fusion algorithm.

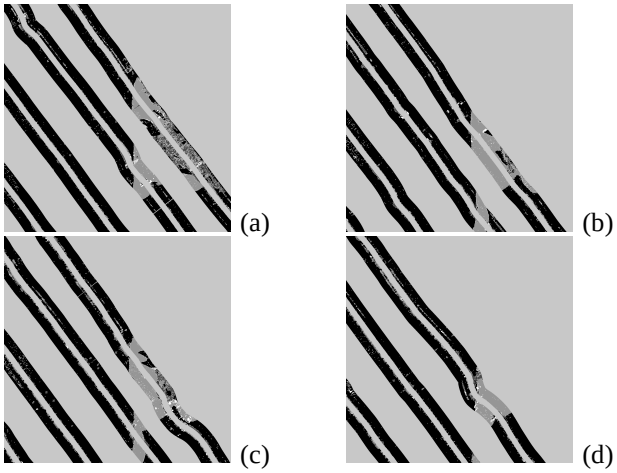

Fig. 13. **(a)-(d) Show four mosaics displaying 15 side-scan tracks of the REMUS Mission .**

The mosaics in Figs. 13 contain a maximum of 5 classes. Each pixel is considered to belong to the sand, ripple, complex, *unmeasured* or *unclassified* class. The large light grey regions are *unmeasured* regions over which the AUV has not passed. The white regions are areas which remain *unclassified* following the geo-referencing process. The geo-referencing process sometimes stretches or contracts pixels from the individual input images to allow mosaics of the correct resolution to be created. The initialization from the Voting Scheme and the Final Fused Result for these mosaics can be seen in Figs. 14 and 15 respectively.

Both the initialization and the final fused results in Figs. 14 to 15 contain much more information than any of the mosaics considered in isolation. The initialization result contains large regions of *unclassified* data where the voting scheme has failed to confidently allocate a seafloor class. These regions have been classified within the final fused result. The regions of seafloor which are classified as *unmeasured* in all the input mosaics have also been left as *unmeasured*. The final output

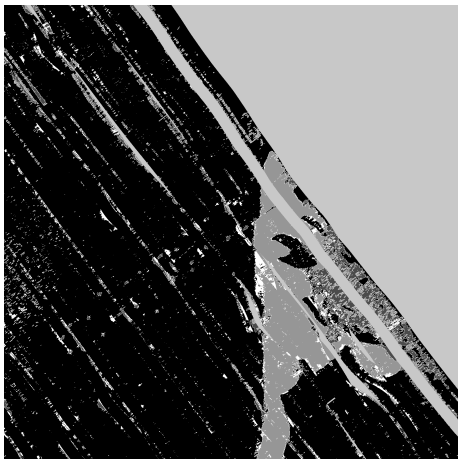

Fig. 14. **The Initialization of the Voting/MRF fusion model after the voting scheme for the Sector 2 mosaics. The white pixels describe *unclassified* regions.**

Fig. 15. **The Final Fused result obtained from the Voting/MRF fusion model for the Sector 2 mosaics. The MRF section of the model has ensured no *unclassified* pixels remain.**

result has fused all the input mosaics to produce a smoothed map where all the pixels where data of some description has been received by the AUV have been classified. The fused map allows a more complete picture of the seafloor to be built up than is possible by considering a single mosaic in isolation.

The Fusion model has produced a more complete and useful picture of the seafloor. The information from the individual mosaics has been fused to produce a smoothed map of the seafloor region. All *unclassified* regions from the initialization result in Fig. 14 have been classified by the Markovian aspect of the model. The benefits of the Fusion model can be clearly seen in the example shown in Figs. 15.

VII. Conclusions

This paper has presented a method for creating and fusing classified sidescan sonar mosaics of the seafloor. An introduction to sidescan sonar imagery was first presented. A normalization and classification model was then detailed which allowed inherent sidescan sonar problems such as the beampattern to be considered. This produced robust classifica-

tion results (segmenting the images into regions of flat seafloor, sand ripples and complex regions) to be obtained for single sidescan images.

A mosaicing algorithm was then presented. This used CML techniques to produce high quality mosaics of the individual classification maps. The generation of these mosaics, where the images are geo-referenced in space, allows the possibility of multi-mosaic fusion.

Two models were then proposed for the fusion of the classified mosaics. The first used voting schemes to initialize the fusion map, after which a Markov model was used to both classify regions previously labeled as *unclassified* and smooth the final result. The second model adopted a probabilistic framework and allowed the reliability of each source to be considered during the fusion process.

The Voting/MRF model was then demonstrated on real classified sidescan mosaics. These mosaics were first produced using the classification and mosaicing models presented in the first section of this paper. Information from the individual mosaics was fused to produce a Fusion map of the entire region of seafloor surveyed by the AUV. The Fusion Model allowed a complete picture of the survey region to be built which would not have been possible from considering any of the mosaics in isolation.

Future research will concentrate on automated methods of obtaining classifier reliability information for the sidescan sonar sensors. If this is available, the probabilistic Fusion model may be implemented. One possibility which will be looked at is whether the result from the MRF/Voting model can be used as an approximation to the ground truth, allowing estimates to the reliability information to be automatically obtained. The mosaicing model will also be enhanced to include existing object detection techniques [69], [70] and automatic data association strategies [71] to simplify and minimize the operators task.

ACKNOWLEDGEMENTS

The authors would like to thank the SACLANT Undersea Research Centre, the US Office of Naval Research and The Woods Hole Oceanographic Institution for allowing the inclusion of data from the BP'02 experiment.

This research work has been partially supported by the 5th Framework Program of research of the European Community through the project AMASON (EVK3-CT-2001-00059).

REFERENCES

[1] 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 *Proc. Autonomous Underwater Vehicle Technology (AUV'94)*, July 1994, pp. 148–155.

[2] A. Ishoy, "How to make survey instruments "AUV-friendly"," in *Proc. MTS/IEEE Internationnal Conference OCEANS'00*, 2000, pp. 1647–1652.

[3] C. von Alt, B. Allen, T. Austin, N. Forrester, R. Goldsborough, M. Purcell, and R. Stokey, "Hunting for mines with REMUS: a high performance, affordable, free swimming underwater robot," in *Proc. MTS/IEEE Internationnal Conference OCEANS'01*, 2001, pp. 117–122.

[4] K. J. Cuevas, M. V. Buchanan, and D. Moss, "Utilizing side scan sonar as an artificial reef management tool," in *Proc. MTS/IEEE Internationnal Conference OCEANS'02*, 2002, pp. 136–140.

[5] D. Diaz, K. Cuevas, M. V. Buchanan, S. Gordon, and W. S. Perret, "Side scan sonar in oyster management," in *Proc. MTS/IEEE Internationnal Conference OCEANS'02*, 2002, pp. 141–145.

[6] B. G. Gonzalez, Y. Petillot, and C. Smith, "Detection and classification of trawling marks in side scan sonar images," in *Proc. Advances in Technology for Underwater Vehicles*, Mar. 2004.

[7] S. Reed, J. Bell, and Y. Petillot, "An unsupervised approach to the detection and extraction of mine features in side scan sonar," in *CAD/CAC 2001 Conference*, Halifax, Canada, Nov. 2001.

[8] R. Urick, *Principles of Underwater Sound*. New York: McGraw-Hill, 1975.

[9] C. Mazel, "Side scan sonar record interpretation," Klein Associates, Inc., New Hampshire, USA, Tech. Rep., 1985.

[10] APL-UW, *APL-UW High-frequency ocean environmental acoustic models handbook*. University of Washington: Applied Physics Laboratory, 1994.

[11] J. M. Bell, "A model for the simulation of sidescan sonar," Ph.D. dissertation, Heriot-Watt University, Edinburgh, Scotland, Sept. 1995.

[12] C. Capus, I. Tena Ruiz, and Y. Petillot, "Compensation for changing beam pattern and residual tvg effects with sonar altitude variation for sidescan mosaicing and classification," in *Proc. 7th. European Conference on Underwater Acoustics*, Delft, The Netherlands, 2004.

[13] S. Anstee, "Removal of range-dependent artifacts from sidescan sonar imagery," DSTO Aeronautical and Maritime Research Laboratory, Australia, Tech. Rep. DSTO-TN-0354, Apr. 2001.

[14] R.D.Muller, N.C.Overkov, J-Y.Royer, A.Dutkiewics, and J.B.Keene, "Seabed classification of the south tasminrise from simrad em12 backscatter data using artificial neural networks," *Australian Journal of Earth Sciences*, pp. 689–700, 1997.

[15] W.K.Stewart, M.Marra, and M.Jiang, "A hierarchical approach to seafloor classification using neural networks," in *Proc. IEEE Int. OCEANS Conf.*, Oct. 1992, pp. 109–113.

[16] B.Bourgeois and C.Walker, "Sidescan sonar image interpretation with neural networks," in *Proc. IEEE Int. OCEANS Conf.*, Honolulu, Hawaia, USA, 1991, pp. 1687–1694.

[17] W. Stewart, M. Jiang, and M. Marra, "A neural network approach to classification of sidescan sonar imagery from a midocean ridge area," *IEEE J. Oceanic Eng.*, no. 2, pp. 214–224, Apr. 1994.

[18] B.Zerr, E.Maillard, and D.Gueriot, "Seafloor classification by neural hybrid system," in *Proc. IEEE Int. OCEANS Conf.*, 1994, pp. 239–243.

[19] D.Alexandrou and D.Pantzartzis, "Seafloor classification with neural networks," in *Proc. IEEE Int. OCEANS Conf.*, Washington D.C, USA, 1990, pp. 18–23.

[20] Z. Reut, N. Pace, and M. Heaton, "Computer classification of seabeds by sonar," *Nature*, pp. 426–428, 1985.

[21] N. Pace and H. Gao, "Swathe seabed classification," *IEEE J. Oceanic Eng.*, no. 2, pp. 83–90, 1988.

[22] D. Tamsett, "Seabed characterisation and classification from the power spectra of sidescan sonar data," *Marine Geophysical Researches*, pp. 43–64, 1993.

[23] D.R.Carmichael, L.M.Linnett, S.J.Clarke, and B.R.Calder, "Seabed classification through multifractal analysis of sidescan sonar imagery," *IEE Radar, Sonar and Navigation*, no. 3, pp. 140–148, 1996.

[24] L.M.Linnett, D.R.Carmichael, S.J.Clarke, and A.D.Tress, "Texture analysis of sidescan sonar data," in *Proc. IEE Conf. Texture Analysis in Radar and Sonar*, London, Nov. 1993.

[25] B. Dasarathy and E. Holder, "Image characterizations based on joint gray-level run-length distributions," *Pattern Recognition Letters*, pp. 497–502, 1991.

[26] R. Haralick, K. Shanmugam, and I. Dinstein, "Textrural features for image classification," *IEEE Trans. Syst., Man, Cybern.*, no. 6, pp. 610–621, Nov. 1973.

[27] S.Subramaniam, H.Barad, and A.B.Martinez, "Seafloor characterization using texture," in *Proc. IEEE South East Conf.*, New Orleans, USA, April 1993.

[28] M.Mignotte, C.Collet, P.Perez, and P.Bouthemy, "Fuzzy logic modelling in sonar imagery: Application to the classification of underwater floor," *Computer Vision and Image Understanding*, pp. 4–24, 2000.

[29] D.T.Cobra and H. Moraes, "Classification of sidescan sonar images through parametric modelling," in *Proc. IEEE Int. OCEANS Conf.*, Brest, France, 1994, pp. 109–113.

[30] D. Gueriot, "Bathymetric and side-scan data fusion for sea-bottom 3D mosaicing," in *Proc. MTS/IEEE Internationnal Conference OCEANS'00*, 2000, pp. 1663–1668.

[31] I. Tena Ruiz, Y. Petillot, and D. M. Lane, "Improved AUV navigation using side-scan sonar," in *Proc. MTS/IEEE Internationnal Conference OCEANS'03*, 2003, pp. 1261–1268.

[32] J. J. Leonard, R. N. Carpenter, and H. J. S. Feder, "Stochastic mapping using forward look sonar," in *Proceedings of International Conference on Field and Service Robotics*, Pittsburgh, Pennsylvania, USA, August 1999, pp. 69–74.

[33] I. Tena Ruiz, Y. Petillot, D. M. Lane, and C. Salson, "Feature extraction and data association for AUV concurent mapping and localisation," in *Proc. IEEE Internationnal Conference on Robotics and Automation (ICRA'01)*, Seoul, Korea, May 2001, pp. 2785–2790.

[34] I. Tena Ruiz, S. de Raucourt, Y. Petillot, and D. M. Lane, "Concurrent mapping and localisation using side-scan sonar," *IEEE J. Oceanic Eng.*, vol. 29, no. 2, pp. 442–456, Apr. 2004.

[35] H.Singh, C.Roman, L.Whitcomb, and D.Yoerger, "Advances in fusion of high resolution underwater optical and acoustic data," in *Proc. IEEE Int. Symposium on Underwater Technology*, May 2000, pp. 206–211.

[36] J.Wright, K.Scott, T-H.Chao, and B.Lau, "Advances in fusion of high resolution underwater optical and acoustic data," in *Proc. IEEE Symposium on Autonomous Underwater Vehicle Technology*, June 1996, pp. 167–175.

[37] A.Solberg, A.Jain, and T.Taxt, "Multisource classification of remotely sensed data: Fusion of landsat TM and SAR images," *IEEE Trans. Geoscience and Remote Sensing*, no. 4, pp. 768–778, JULY 1994.

[38] I.Bloch, "Information combination operators for data fusion: A comparative review with classification," *IEEE Trans. Systems, Man and Cybernetics-Part A: Systems and Humans*, no. 1, pp. 52–67, JAN 1996.

[39] ——, "Some aspects of Dempster-Shafer evidence theory for classification of multi-modality medical images taking partial volume into account," *Pattern Recognition Letters*, pp. 905–919, 1996.

[40] S.Foucher, M.Germain, J-M.Boucher, and G.Benie, "Multisource classification using ICM and Dempster-Shafer theory," *IEEE Trans. Instrumentation and Measurement*, no. 2, pp. 277–281, APRIL 2002.

[41] L.Fouque, A.Appriou, and W.Pieczynski, "An Evidential Markovian model for data fusion and unsupervised image classification," in *Proc. 3rd ISIF Int. Conf. Information Fusion*, 2000, pp. 25–32.

[42] J.Kittler, M.Hatef, R.Duin, and J.Matas, "On combining classifiers," *IEEE Trans. Pattern Analysis and Machine Intelligience*, no. 3, pp. 226–239, MARCH 1998.

[43] L.Xu, A.Krzyzak, and C.Y.Suen, "Methods of combining multiple classifiers and their applications to handwriting recognition," *IEEE Trans. Systems,Man and Cybernetics*, no. 3, pp. 418–435, MAY/JUNE 1992.

[44] S. Hegarat-Mascle, I.Bloch, and D.Vidal-Madjar, "Application of dempster-shafer evidence theory to unsupervised classification in multisource remote sensing," *IEEE Trans. Geoscience and Remote Sensing*, no. 4, pp. 1018–1031, JULY 1997.

[45] J. Cleynenbreugel, S.A.Osinga, F.Fierens, P.Suetens, and A.Oosterlinck, "Road extraction from multi-temporal satellite images by an evidential reasoning approach," *Pattern Recognition Letters*, pp. 371–380, JUNE 1991.

[46] L.Fouque, A.Appriou, and W.Pieczynski, "Multiresolution hidden Markov chain model and unsupervised image segmentation," in *Proc. 4th IEEE Southwest Symposium on Image Analysis and Interpretation*, 2000, pp. 121–125.

[47] A.Bendjebbour, Y.Delignon, L.Fouque, V.Samsom, and W.Pieczynski, "Multisensor image segmentation using Dempster-Shafer fusion in Markov fields context," *IEEE Trans. Geoscience and Remote Sensing*, no. 8, pp. 1789–1798, AUGUST 2001.

[48] S. Hegarat-Mascle, A.Quesney, D.Vidal-Madjar, and O.Taconet, "Land cover discrimination from multitemporal ERS images and multispectral landsat images: A study case in an argicultural area in france," *Int. Journal Remote Sensing*, no. 3, pp. 435–456, 2000.

[49] W.A.Wright, "Fast image fusion with a markov random field," in *Proc. IEE Conf. Image Processing and its Applications*, 1999, pp. 557–561.

[50] D.Geman, S.Geman, C.Graffigne, and P.Dong, "Boundary detection by constrained optimization," *IEEE Trans. Pattern Analysis and Machine Intelligience*, no. 7, pp. 609–628, JULY 1990.

[51] J.Besag, "On the statistical analysis of dirty pictures," *J. R. Statist. Soc.*, no. 3, pp. 261–279, 1986.

[52] J.A.Benediktsson and I.Kanellopoulos, "Classification of multisource and hyperspectral data based on decision fusion," *IEEE Trans. Geoscience and Remote Sensing*, no. 3, pp. 1367–1377, MAY 1999.

[53] N. A. Brokloff, "Matrix algorithm for doppler sonar navigation," in *Proc. IEEE Internationnal Conference OCEANS'94*, 1994, pp. III/378 – III/383.

[54] R. Cox and S. Wei, "Advances in the state of the art for AUV inertial sensors and navigation systems," *IEEE J. Oceanic Eng.*, vol. 20, no. 4, pp. 361–366, Oct. 1995.

[55] G. Grenon, P. E. An, S. M. Smith, and A. J. Healey, "Enhancement of the inertial navigation system for the Morpheus autonomous underwater vehicles," *IEEE J. Oceanic Eng.*, vol. 26, no. 4, pp. 548–560, Oct. 2001.

[56] K. Vickery, "Acoustic positioning systems. a practical overview of current systems," in *Proc. Autonomous Underwater Vehicle Technology (AUV'98)*, Aug. 1998, pp. 5–17.

[57] O. Bergem, "Bathymetric navigation of autonomous underwater vehicles using a multibeam sonar and a Kalman filter with relative measurement covariance matrices," Ph.D. dissertation, Univeristy of Trondheim, Norway, Dec. 1993.

[58] M. Sistiaga, J. Opderbecke, and M. Aldon, "Depth image matching for underwater vehicle navigation," in *International Conference on Image Analysis and Processing*, 1999, pp. 624 –629.

[59] E. Coiras, I. Tena Ruiz, Y. Petillot, and D. M. Lane, "Fusion of multiple side-scan sonar views," in *Proc. IEEE Internationnal Conference OCEANS'04*, 2004, pp. ?–?

[60] P. Maybeck, *Stochastic models, estimation ans control. Volume 2*, ser. Mathematics in Science and Engineering. Academic Press, 1982, vol. 141-2.

[61] Y. Bar-Shalom and T. Fortmann, *Tracking and data association.*, ser. Mathematics in Science and Engineering. Academic Press, 1988, vol. 179.

[62] 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.

[63] M. W. M. G. Dissanayake, H. Durrant-Whyte, S. Clark, and M. Csorba, "A solution to the simultaneous localisation and map building (SLAM) problem," the University of Sydney, Sydney, Australia, Tech. Rep. ACFR-TR-01-99, Jan. 1999.

[64] A. Gelb, Ed., *Applied Optimal Estimation*. Cambridge, MA, USA: The M.I.T. Press, 1974.

[65] M.Mignotte, C.Collet, P.Perez, and P.Bouthemy, "Sonar image segmentation using an unsupervised hierarchical MRF model," *IEEE Trans. Image Processing*, no. 7, pp. 1216–1231, JULY 2000.

[66] H. Derin and H. Elliot, "Modeling and segmentation of noisy and textured images using Gibbs random field," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 9, no. 1, pp. 39–55, 1987.

[67] MeasTex, "MeasTex Image Texture Database and Test Suite," available online at: http://www.cssip.uq.edu.au/meastex/meastex.html, May 1997.

[68] I. Tena Ruiz, Y. Petillot, D. M. Lane, and A. Cormack, "Large scale side-scan sonar mosaics." in *Proc. Maritime Reconnaissance for NATO's Recognized Environmental Picture (MREP)*, La Spezia, Italy, May 2003.

[69] S. Reed, J. Bell, and Y. Petillot, "Unsupervised segmentation of object shadow and highlight using statistical snakes," in *GOATS 2000 Conference*, La Spezia, Italy, Aug. 2001.

[70] D. M. Lane, M. J. Chantler, and D. Y. Dai, "Robust tracking of multiple objects in sector scan sonar image sequences using optical flow motion estimation," *IEEE J. Oceanic Eng.*, vol. 23, no. 1, pp. 31–46, Jan. 1998.

[71] I. Tena Ruiz, "Enhanced concurrent mapping and localisation using forward-looking sonar," Ph.D. dissertation, Heriot Watt University, Edinburgh, Scotland, Sept. 2001.

Scott Reed Scott Reed graduated from Edinburgh University in 1999 with a MPhys Astrophysics (Hons). He then obtained an MSc and his PhD in the Ocean Systems Lab at Heriot-Watt University in 2000 and 2004 respectively. His PhD thesis examined automated detection, classification and segmentation techniques for use with side-scan sonar imagery. He is currently working at SeeByte Ltd in Edinburgh, developing Data Fusion techniques for the UK MoD BAUUV project. His research interests include Markov models, Statistical Image Processing, Texture analysis and Data Fusion.

Ioseba Tena Ruiz received a B.Eng (Hons) in Electronic Engineering from Heriot-Watt University in 1996 and his PhD in 2001 from the same university. He has worked as a Research Associate at Heriot-Watt University in five different research projects CLASS (CMPT), ARAMIS (EU), AUTOTRACKER (EU), ALIVE (EU) and AMASON (EU). His interests include navigation systems, concurrent mapping and localisation, processing sonar returns, classifying sonar returns and obstacle avoidance systems for autonomous underwater-vehicles. He is currently working on simultaneous mapping and localisation accross multiple platforms. He is a reviewer for IEE and IEEE journals.

Yvan Petillot gained his PhD from University de Brewage Occidental in real-time pattern recognition using optical processors. He also has a M.Sc. in optics and signal processing and an Engineering degree in Telecommunications with a specialization in Image and signal processing. He is currently lecturer at Heriot Watt University carrying out research and teaching activities in image processing including texture and shape analysis, segmentation and classification of sonar and video images, simultaneous mapping and localisation and navigation. He also has a particular interest in mine dectection and classification using image based statistical techniques. He is the principal investigator on 2 EPSRC Research projects (GR/S68088/01 and GR/S16980/01) on weapons detection using mm-wave imaging (222K) and on multimodal sensor fusion for seabed classification (129K) respectively. He is also co-investigator of the Framework V AUTOTRACKER project (obstacle avoidance and acoustic pipe tracking using an AUV) and the Framework V Amason project (Data fusion of sonar data products for seabed classification and habitat mapping). Each project has a University budget of around 300K. He is also a reviewer for IEE and IEEE journals.