AUTOMATIC SHIP HULL INSPECTION - THE DETECTION OF MINE-LIKE TARGETS IN SONAR DATA USING MULTI-CAD FUSION AND TRACKING TECHNOLOGIES

P.Y. Mignotte SeeByte Ltd., Edinburgh, UK
Dr. S. Reed SeeByte Ltd., Edinburgh, UK
A. Cormack Seebyte Ltd., Edinburgh, UK
Dr. Y. R. Petillot Seebyte Ltd., Edinburgh, UK

1 INTRODUCTION

Ship hull inspection operations are presently carried out manually using divers who are tasked with detecting possible mine-like targets on the sides of the ship. This task is both time consuming and dangerous due to the complex environment and poor water visibility. A tethered platform such as a ROV cannot be deployed in these cluttered areas (pipes, propeller). Autonomous underwater vehicles (AUV's) are now being used to automate this process. These vehicles are capable of surveying the entire ship hull[1,2] solely relying on relative navigation. An embedded computer aided detection and classification (CAD/CAC) system is required to further automate the inspection process. Automatic inspection techniques (such as using the sonar data) are required for complex hull regions where more traditional approaches (such as using a DVL[1] to navigate) would struggle.

The development of CAD/CAC systems has been focussed principally for side-scan sonar(SSS)[3]. These studies have highlighted the difficulty in finding one single model solution to the CAD/CAC problem, which is robust to changes in the sonar, environment and image quality. Multi-model fusion[4] has been demonstrated as a possible solution where the fusion mechanism relieves any individual CAD/CAC model of the responsibility of detecting all the mine-like targets. Fusion also shows potential in identifying and removing false alarm detects. These are unavoidable for any automated system and are especially problematic in cluttered environments such as ship hulls. The success and applicability of a CAD/CAC system will often be dictated by the level of false alarms picked up by the system. A multi-model fusion CAD/CAC system, similar to those employed in SSS systems, is presented in this paper for use on forward looking sonar systems.

This paper demonstrates an unsupervised, real time automated target recognition (ATR) system for hull inspection. Multiple and diverse CAD models process the sonar data after which a fusion model is used to streamline the results. The diversity of the CAD strategies allows a high detection rate while the pruning phase during the fusion process reduces the risk of false alarms. The high frame rate of the hull data (ensuring any object appears in multiple adjacent frames) also allows tracking technologies to also to be used in the CAD/CAC solution. Using the vehicle navigation information, any detection is required to be tracked over a number of consecutive frames before being confirmed as a real object.

This paper is organised as follows. The next section gives an overview of the system. The CAD/CAC and Tracking solutions are described in the third section. Finally, results are shown on data acquired during HullFest'06. Conclusions regarding the possibilities and further refinements are presented in the last section.

2 SYSTEM OVERVIEW

Figure 1 presents a schematic view of the overall system with an emphasis on the CAD/CAC process. Each frame is concurrently processed by n elementary CAD systems (typically n=2, 3). These detectors perform a quick image segmentation into three classes: dark, medium and bright areas. Bright regions often represent regions of high reflectivity such as man-made object while dark regions may represent shadow regions which can also be indicative of a possible object. A run-length algorithm permits labelling of the different regions and the extraction of their features (e.g. width, length). The result of this CAD stage is a list of regions of interest, classified as shadow, background or highlight. Further processing is performed on these regions and their features independent of the image data to ensuring a low computational cost and real time analysis.

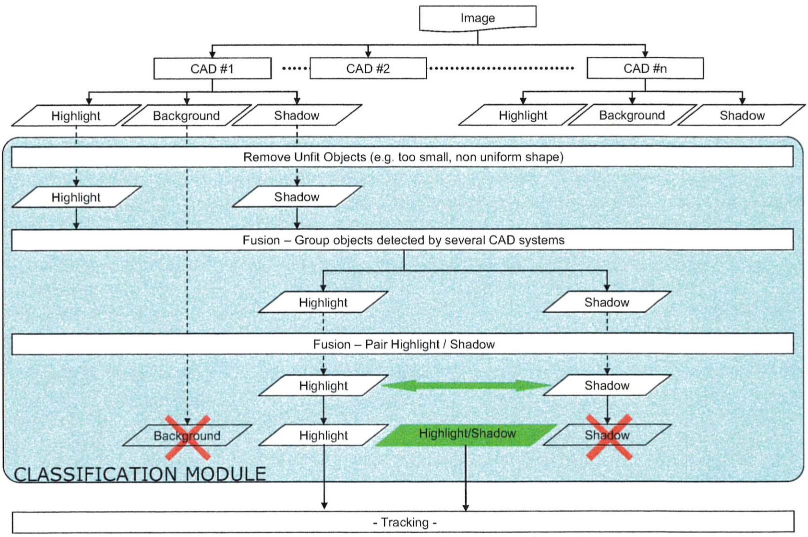

Figure 1: Automated Target Detection Overview

A first pruning is performed to remove objects which are obvious false alarms based on user-defined size restrictions. The detections obtained from the n detectors are then combined to remove redundant detections picked up by the different CAD's. Redundant features are defined as objects which have been picked up by multiple CAD models. A set of unique background, shadow and highlight objects are obtained and passed onto a pairing module. This module is responsible for automatically associating highlight detection with their cast shadow. This association is sound for the observation of 3D objects lying on a surface. A final pruning is performed to remove detections of low confidence (e.g. only detected by one CAD) or objects which belong to a class of no interest.

The CAD/CAC system is performed on each frame. The final list of relevant objects is passed to the tracking system. This reads in the navigation from the vehicle to estimate the expected motion of each detection between consecutive frames. Consistent detections over a series of frame (typically 5) are validated as objects and output for further actions.

3 TARGET DETECTION AND TRACKING

3.1 Suite of elementary detectors

Four elementary and diverse detectors have been implemented. Their purpose is to quickly provide a list of possible discrepancies within the images. Each discrepancy is a possible threat which must be validated or discarded at the end of the fusion and tracking process. The discrepancy detection is performed based on a segmentation stage and run-length labelling algorithm[7]. Each region receives a unique label and its features are measured. These features are position, class, dimension and area. These are the sole pieces of information retained for subsequent analysis to permit real-time processing.

The first CAD methodology relies on a K-means clustering algorithm[5]. The unsupervised model segments the image into three classes. This raw segmentation is prone to noise as only the greylevel pixel value is considered during the initial segmentation. A simple Markov random field[6] is used to smooth the obtained classmap by reintroducing neighbourhood information.

The second CAD uses a similar approach but is tailored for sonar processing. Two percentages are set and reflect the ratio of pixels per image which are assumed to lie within the dark and bright classes. Usually, echoes are small and localised so only a small percentage of pixels of the image is generally allocated to the bright class. The Markov smoothing stage is applied to ensure a smooth segmentation.

The third CAD considers a region-growing algorithm around the brightest pixels within the image. Two parameters are set: one to trigger an object growing region and another to define how the region growing process is stopped. This method is region based and does not require further smoothing of the class map.

The fourth CAD detects discrepancies using the Student T-Test[8]. Areas where the first and second order statistics deviate from their neighbourhood are highlighted as possible objects.

These four detectors rely on slightly different hypothesis during the segmentation stage. Different false alarms are detected by each model (allowing them to be removed) while it is hoped that real objects are detected and confirmed by multiple CAD models. Section 3.2 describes the fusion of these detections in order to reveal objects and remove false alarms.

3.2 Fusion rule for Redundant Detection grouping and Highlight/Shadow pairing

This section mainly refers to the classification module shown in Figure 1. Before the prospective detections are provided to the tracker, two fusion processes are carried out reduce the number of detected objects. Multiple detections relating to the same object are first merged. The final detection is reduced to the smallest common area of the set of detections. The combination relies on an overlap argument.

During the second fusion phase, joint detections of a bright (object) and a dark (cast shadow) area are merged into a single object. The detection of both a highlight and shadow region increases the confidence of an object being present. The association uses the feature pairing process proposed by Scott and Longuet-Higgins[9]. A position based feature similarity is defined and the pairing estimation is performed using a Singular Value Decomposition. The likelihood, using the standard notation G_{ij}[9,10], of two features to be associated follows a Gaussian model (1):

$$G_{ij} = \cos(\vec{x}_H - \vec{x}_S, \vec{\mu}) \cdot \exp\left[(\vec{x}_H - \vec{x}_S - \vec{\mu}) \cdot \Sigma^{-1} \cdot (\vec{x}_H - \vec{x}_S - \vec{\mu})^T\right] \tag{1}$$

Where: X_H and X_S correspond to the position of a highlight and shadow.

μ is an estimation of the ideal relative position between an object and its cast shadow.

Σ is the covariance matrix denoting the possible error between the ideal and real relative position of a highlight and its shadow.

The model is constrained by the description of the required orientation between the highlight and shadow regions. The cosine term prevents a cast shadow appearing up-range from an object. The covariance Σ allows a greater error in the direction of observation as the localisation of the shadow is less accurate.

3.3 Vehicle navigation based tracking

At the end of the fusion process, a short list of relevant objects is produced. This contains a list of alarms detected on a single frame. These alarms are entered into the tracking module which follows an object over a number of consecutive frames before confirming a detection as an object.

For the sake of simplicity, our tracking method uses a model based approach. The vehicle provides a hull relative navigation based on a Doppler velocity log (DVL)[1]. This allows estimation of how much the frame is modified due to the displacement of the vehicle.

Detections in the previous frame are displaced according to this estimate. An association based on the size and position of the objects is performed between the previous and current frame. The score of the detections depends on their frequency of detection. A tracked detection is usually validated if it has been detected in at least 70% of the frames, over more than 5 frames.

4 RESULTS

4.1 Experimental Platform

The results presented in this paper were collected using the hover capable HAUV vehicle developed by Bluefin Robotics[1]. The HAUV vehicle uses a Doppler Velocity Log (DVL) to provide a hull relative navigation. The vehicle follows a lawn-mover type trajectory at about one meter from the hull. The observation is performed using a Didson sonar camera. This sonar provides high resolution imagery (~7cm) of the ship hull.

4.2 Ship Hull Inspection

This section presents results of the detection process using different datasets courtesy of Bluefin Robotics. These datasets consist of a Didson data file and HAUV navigation file from surveys which were around 6 minutes long. This data was gathered during the US Hullsfest '06 trials.

The whole CAD/CAC process is running at the sonar frame rate (5 frames per second) ensuring the system may run embedded and real time. The first 2 CAD models from the suite of CADs were used in this application.

Each screenshot presents a Didson frame in a Cartesian coordinates system. Grey-level rectangles represent the areas detected as possible objects at each frame. The grey level of the line increases as the confidence the detection is actually an object increases over subsequent frames. Thick yellow boxes highlight picked up objects. The bottom line, if displayed, defines the line above which is assumed to contain hull information. The presence of the hull is also automatically detected and will later be used to servo the vehicle automatically off the hull using only the sonar data.

Figure 2 presents 6 consecutive frames of didson data. The vehicle is flying to the left and an object can be seen first appearing in frame 818. This object is consistently detected over the following frames and its confidence increases up it is finally confirmed as an object. In this example, the highlight and cast shadow have been paired together as a single object, further increasing the confidence.

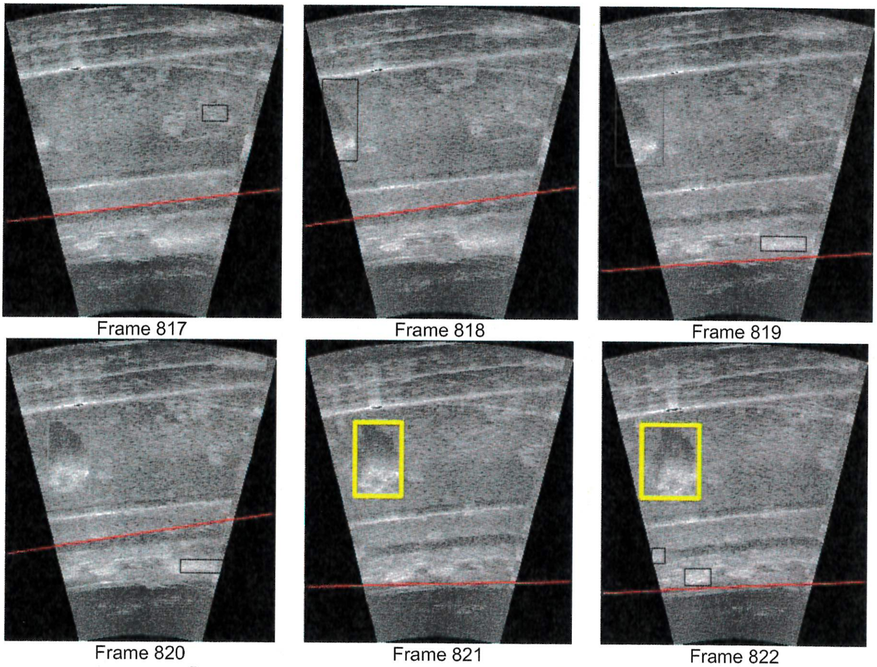

Figure 2: Tracking and Highlight/Shadow pairing capabilities

Figure 3 presents another series of frames extracted during a survey containing a very cluttered area. The images contain pipes which are located under the ship and tend to generate a high number of false alarms. However, all false detections are discarded during the tracking process until the object of frame 751 is eventually detected. Typically, this frame presents a lot of false detection which do not trigger a false alarm thanks to the tracking capabilities.

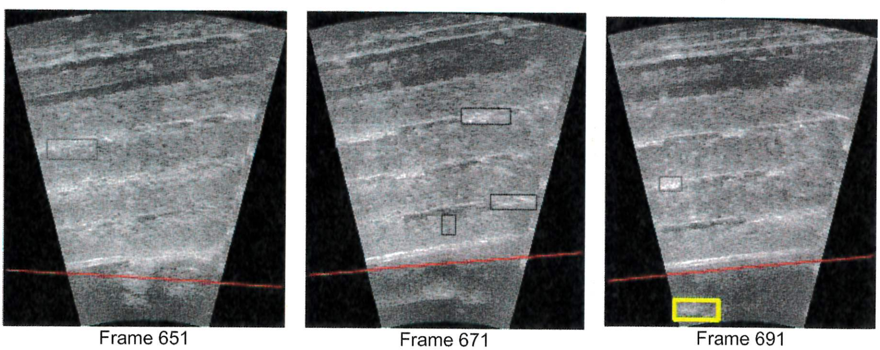

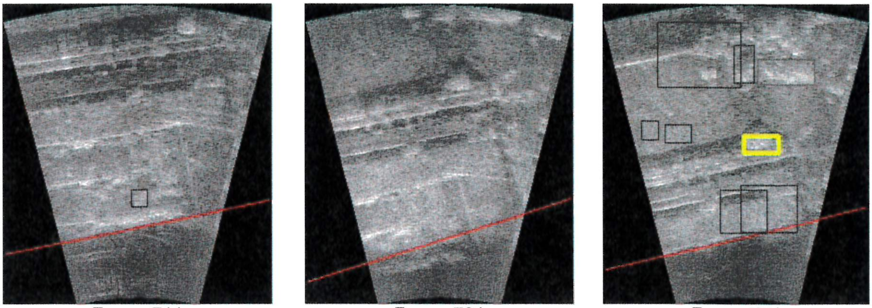

Frame 711 Frame 731 Frame 751

Figure 3: CAD/CAC performance over a cluttered area

4.3 Mine Neutralization Feature Based Navigation Applications

The concurrent use of several simple detectors and a general fusion scheme achieves a high robustness which can be applied to other ATR applications. This scheme has been successfully applied to the Mine Neutralization problem and will be used to automatically provide features for a Feature Based Navigation (FBN) system developed by MIT. A Nekton ranger vehicle mounted with a BlueView blazed array sonar is flown through a field of sonar reflectors. The accurate detection of the sonar reflectors is crucial to refine the navigation estimate (using a Simultaneous Localisation and Mapping (SLAM) approach) to allow the AUV to home in on the mine to be neutralized. The system described in this paper has been applied to this problem and has been shown to be robust at accurately detecting the features within the simulated mine field.

Figure 4 presents a sequence of 6 frames containing BlueView data. This dataset has been acquired within the ONR funded Mine Neutralisation project. The fusion relies on the use of the first and third CAD system of our suite. The unanimous vote rule discards all the noise due to the sensor (appearing in the central and side areas of the sonar field of view). These regions are a common source of false alarms with classical methods (around 10 false detections per frame). No false alarms are picked up while all real features have been detected and tracked.

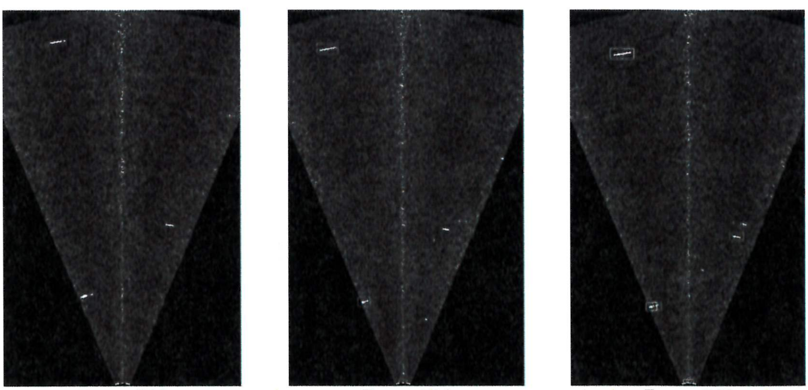

Frame 730 Frame 731 Frame 732

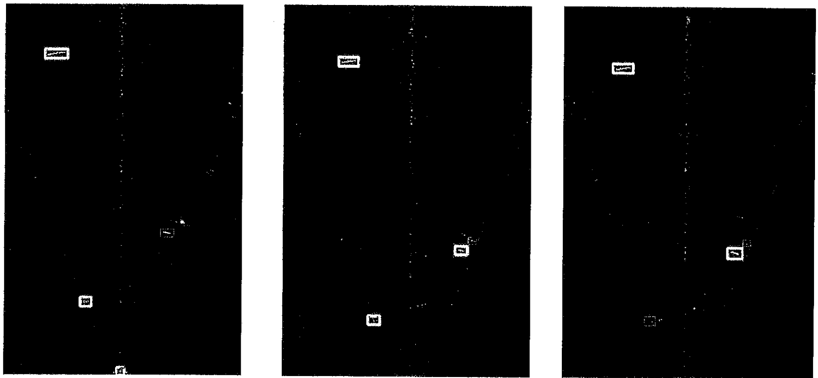

Frame 733 Frame 734 Frame 735

Figure 4: Feature detections for SLAM based navigation

5 CONCLUSION

This paper has introduced an automated target recognition system. Based on a suite of simple CAD, a wide range of objects can be detected independently of the limitations of a single system. A fusion stage permits the set of detections to be fused and pruned in order to obtain a reduced list of individual objects. The latter are fed into a model-based tracking system which relies on the vehicle navigation to predict the position of the detections from one frame to the other. This process is used to confirm real threats and remove false alarms.

The system has been tested on a wide range of datasets showing its robustness and efficiency. A high detection rate with a low number of false alarms has been achieved even with a low computational load. Additionally, the proposed approach is versatile and has been successfully applied to different applications where an embedded CAD/CAC is required.

Further work will aim at adding new simple detectors to increase the variability of the detection. Opportunities for machine-learning based fusion are being considered to enhance the current model-based system, which requires some manual parameterisation. Ongoing work focuses on the performance of particle filter models for simultaneous fusion and tracking.

Acknowledgement:

The authors would to thank Dr. Jerome Vaganay(Bluefin Robotics) and Dr. Lee Thompson(Blueview Technologies) for providing the datasets used to assess this system. We would like to acknowledge support from the Office of Naval Research under the contract number 12-458-456-729.

6 REFERENCES

1. J. Vaganay, M.L. Elkins, S. Willcox, F.S. Hover, R.S. Damus, S. Desset, J.P. Morash and V.C. Polidoro, Ship Hull Inspection by Hull-Relative Navigation and Control, Proceedings of MTS/IEEE Oceans 2005, Vol.1, 761—766, 2005.
2. S. Reed, A. Cormack, K. Hamilton, I. Tena Ruiz and D.Lane, Automatic Ship Hull Inspection using Unmanned Underwater Vehicles, Proceedings of the 7th Int. Symposium on Technology and the Mine Problem, Naval Postgraduate School, Monterey Bay, USA, 2006.
3. S. Reed, Y.R. Pétillot and J. Bell, A model based approach to mine detection and classification in sidescan sonar, Proceedings of Oceans 2003, Vol. 3, 1402—1407, 2003.
4. G.J. Dobeck, Fusing sonar images for mine detection and classification, In Proceedings of SPIE, Vol. 3710, 602—614, 1999.
5. R.O. Duda, P.E. Hart and D.G. Stork, Pattern Classification, 2nd edition, Wiley Interscience, 2000.
6. S. Geman and D. Geman, Stochastic relaxation, Gibbs distribution and the Bayesian restoration of images, IEEE Trans. on PAMI, 6(6):721—741, 1984.
7. R. M. Haralick and L. G. Shapiro, Computer and Robot Vision, Volume I, Addison-Wesley, 1992.
8. W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, Numerical Recipes in C++, Cambridge University Press, 2nd edition, 2002.
9. G. Scott and H. Longuet-Higgins, An algorithm for associating the features of two patterns. Proceedings of Royal Society of London, B244:21-- 26, 1991.
10. M. Pilu and A. Lorusso, Uncalibrated Stereo Correspondence based on Singular Value Decomposition, In Proceedings of BMVC, 1997.